

Guide to Creating VMS Modular Procedures

Order Number: AA-LA58A-TE

April 1988

This document describes how to design and code procedures that conform to the VMS Modular Programming Standard. It also describes how to insert modules in an object module library, shareable image, or shareable image library.

Revision/Update Information: This manual supersedes the *Guide to Creating Modular Procedures on VAX/VMS, Version 4.2.*

Software Version: VMS Version 5.0

**digital equipment corporation
maynard, massachusetts**

April 1988

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license and may be used or copied only in accordance with the terms of such license.

No responsibility is assumed for the use or reliability of software on equipment that is not supplied by Digital Equipment Corporation or its affiliated companies.

Copyright ©1988 by Digital Equipment Corporation

All Rights Reserved.
Printed in U.S.A.

The postpaid READER'S COMMENTS form on the last page of this document requests the user's critical evaluation to assist in preparing future documentation.

The following are trademarks of Digital Equipment Corporation:

DEC	DIBOL	UNIBUS
DEC/CMS	EduSystem	VAX
DEC/MMS	IAS	VAXcluster
DECnet	MASSBUS	VMS
DECsystem-10	PDP	VT
DECSYSTEM-20	PDT	
DECUS	RSTS	
DECwriter	RSX	

digital™

ZK4518

**HOW TO ORDER ADDITIONAL DOCUMENTATION
DIRECT MAIL ORDERS**

USA & PUERTO RICO*

Digital Equipment Corporation
P.O. Box CS2008
Nashua, New Hampshire
03061

CANADA

Digital Equipment
of Canada Ltd.
100 Herzberg Road
Kanata, Ontario K2K 2A6
Attn: Direct Order Desk

INTERNATIONAL

Digital Equipment Corporation
PSG Business Manager
c/o Digital's local subsidiary
or approved distributor

In Continental USA and Puerto Rico call 800-258-1710.

In New Hampshire, Alaska, and Hawaii call 603-884-6660.

In Canada call 800-267-6215.

* Any prepaid order from Puerto Rico must be placed with the local Digital subsidiary (809-754-7575).

Internal orders should be placed through the Software Distribution Center (SDC), Digital Equipment Corporation, Westminister, Massachusetts 01473.

Production Note

This book was produced with the VAX DOCUMENT electronic publishing system, a software tool developed and sold by DIGITAL. In this system, writers use an ASCII text editor to create source files containing text and English-like code; this code labels the structural elements of the document, such as chapters, paragraphs, and tables. The VAX DOCUMENT software, which runs on the VMS operating system, interprets the code to format the text, generate a table of contents and index, and paginate the entire document. Writers can print the document on the terminal or line printer, or they can use DIGITAL-supported devices, such as the LN03 laser printer and PostScript[™] printers (PrintServer 40 or LN03R ScriptPrinter), to produce a typeset-quality copy containing integrated graphics.

[™] PostScript is a trademark of Adobe Systems, Inc.

Introduction

The purpose of this study is to investigate the effects of the proposed system on the performance of the system. The study is divided into two main parts: a theoretical analysis and an experimental evaluation. The theoretical analysis is based on the principles of the system and the results of previous studies. The experimental evaluation is based on the results of a series of experiments conducted on a real system. The results of the study are presented in the following sections.

Contents

PREFACE

xiii

CHAPTER 1 INTRODUCTION TO MODULAR PROCEDURES

1-1

1.1 FOLLOWING THE SOFTWARE LIFE CYCLE

1-1

1.2 DEFINING THE MODULAR PROCEDURE

1-1

1.3 EXISTING VMS SYSTEM PROCEDURES

1-4

1.3.1 Run-Time Library Procedures

1-6

1.3.2 System Services

1-8

1.3.3 Utility Routines

1-10

1.3.4 Record Management Services

1-11

1.4 TOOLS FOR DEVELOPING APPLICATIONS ON VMS

1-12

CHAPTER 2 DESIGN

2-1

2.1 ORGANIZING NEW APPLICATIONS

2-1

2.1.1 Organizing Files and Modules

2-1

2.1.2 Organizing Procedures into Modules

2-1

2.2 DEFINING A MODULAR PROCEDURE INTERFACE

2-3

2.2.1 Explicit Arguments

2-3

2.2.2 Implicit Arguments

2-3

2.2.2.1 Implicit Arguments Allocated by the Calling Program • 2-5

2.2.2.2 Implicit Arguments Allocated by the Called Procedure • 2-5

2.2.3 How to Avoid Using Implicit Arguments

2-5

2.2.3.1 Combining Procedures • 2-6

2.2.3.2 User Action Routine • 2-7

2.2.3.3 Designating Responsibility to the Calling Program • 2-8

2.2.4 Order of Arguments

2-11

2.2.5 Using Optional Arguments

2-11

2.3 JSB ENTRY POINTS

2-12

Contents

2.4	USING SYSTEM RESOURCES	2-12
2.4.1	Choosing a Storage Type	2-12
2.4.1.1	Stack Storage • 2-12	
2.4.1.2	Heap Storage • 2-12	
2.4.1.3	Static Storage • 2-13	
2.4.1.4	Avoiding Use of Static Storage • 2-14	
2.4.1.5	Summary of Storage Use by Language • 2-15	
2.4.2	Using Event Flags	2-16
2.4.3	Using Logical Unit Numbers	2-16
2.5	USING INPUT/OUTPUT	2-16
2.5.1	Terminal Input/Output	2-17
2.5.2	File Input/Output	2-18
2.6	BEGINNING THE INTERNAL DOCUMENTATION	2-19
2.6.1	How to Write a Module Description	2-19
2.6.2	How to Write a Procedure Description	2-20
2.7	PLANNING FOR SIGNALING AND CONDITION HANDLING	2-22
2.7.1	Guidelines for Signaling Error Conditions	2-23
2.7.2	Guidelines for Returning Condition Values	2-23
2.7.3	When to Signal or Return Condition Values	2-24
CHAPTER 3 CODING		3-1
3.1	CODING GUIDELINES	3-1
3.1.1	Writing Position-Independent Code (Required)	3-1
3.1.2	Adhering to the Naming Conventions	3-1
3.1.2.1	Facility Naming Conventions (Recommended) • 3-2	
3.1.2.2	Procedure Naming Conventions (Recommended) • 3-3	
3.1.2.3	File Naming Conventions (Recommended) • 3-4	
3.1.2.4	Module Naming Conventions (Required) • 3-4	
3.1.2.5	PSECT Naming Conventions (Required) • 3-5	
3.1.2.6	Lock Resource Naming Conventions (Recommended) • 3-5	
3.1.2.7	Global Variable Naming Conventions (Recommended) • 3-5	
3.1.2.8	Status Code and Condition Value Naming Conventions (Required) • 3-6	
3.1.3	Using Common Source Files (Recommended)	3-7
3.1.4	Maintaining Code Readability	3-7
3.1.4.1	Using Symbols in Place of Numbers (Recommended) • 3-8	
3.1.4.2	Using Uppercase and Lowercase Characters (Recommended) • 3-8	
3.1.4.3	Adding Optional Spaces (Recommended) • 3-8	

3.1.4.4	Inserting Block Comments (Recommended) • 3-9	
3.1.5	Using VMS System Services	3-11
3.1.6	Invoking Optional User Action Routines	3-11
3.1.6.1	Procedure Entry Mask • 3-11	
3.1.6.2	Bound Procedure Value • 3-12	
3.2	INITIALIZING MODULAR PROCEDURES	3-12
3.2.1	Initializing Storage	3-14
3.2.2	Testing and Setting a First-Time Flag	3-14
3.2.3	Using LIB\$INITIALIZE	3-17
3.3	WRITING AST-REENTRANT CODE	3-19
3.3.1	What is an AST?	3-19
3.3.2	AST-Reentrancy Versus Full-Reentrancy	3-20
3.3.3	Guidelines for Writing AST-Reentrant Modular Procedures	3-20
3.3.4	How to Eliminate Race Conditions During Concurrent Access	3-21
3.3.4.1	Performing All Accesses in One Instruction • 3-22	
3.3.4.2	Using "Test and Set" Instructions • 3-23	
3.3.4.3	Keeping a Call-in-Progress Count • 3-24	
3.3.4.4	Disabling AST Interrupts • 3-24	
3.3.5	Performing Input/Output at AST-Level	3-25
3.3.6	Condition Handling at AST Level	3-26
CHAPTER 4 TESTING		4-1
4.1	UNIT TESTING	4-1
4.1.1	Black Box Testing	4-2
4.1.2	White Box Testing	4-3
4.2	LANGUAGE-INDEPENDENCE TESTING	4-4
4.3	INTEGRATION TESTING	4-5
4.3.1	The "All at Once" Approach to Integration Testing	4-5
4.3.2	The Incremental Approach to Integration Testing	4-6
4.4	TESTING FOR REENTRANCY	4-6
4.4.1	Checking for AST Reentrancy	4-7
4.4.1.1	Checking for AST Reentrancy using the Debugger • 4-7	
4.4.1.2	Checking for AST Reentrancy by Desk Checking • 4-7	
4.4.2	Checking for Full Reentrancy	4-8

Contents

4.5	PERFORMANCE ANALYSIS	4-8
4.5.1	SHOW Entry Point	4-8
4.5.2	STAT Entry Point	4-9
4.6	MONITORING PROCEDURES IN THE RUN-TIME LIBRARY	4-9
CHAPTER 5 INTEGRATION		5-1
5.1	GROUPING PROCEDURES	5-1
5.1.1	Creating Facility Prefixes	5-1
5.1.2	Creating Object Module Libraries	5-2
5.2	SHAREABLE LIBRARY IMAGES	5-3
5.2.1	Creating Shareable Library Images	5-4
5.2.2	Creating the Transfer Vector	5-5
5.2.3	Creating the Linker Options File	5-8
5.2.4	Creating the Shareable Library Image	5-10
5.2.5	Combining Shareable Images into a Shareable Image Library	5-11
5.3	LINKING TO LIBRARIES OF MODULAR PROCEDURES	5-11
CHAPTER 6 MAINTENANCE		6-1
6.1	UPWARD COMPATIBILITY	6-1
6.1.1	Making Your Procedures Upwardly Compatible	6-1
6.1.2	Regression Testing	6-1
6.1.2.1	Updating the Transfer Vector • 6-3	
6.2	ADDING ARGUMENTS TO EXISTING ROUTINES	6-3
6.2.1	Adding New Arguments to the Procedure	6-4
6.2.2	Using Argument Blocks	6-4
6.3	UPDATING LIBRARIES	6-5
6.3.1	Updating Object Libraries	6-5
6.3.2	Updating Shareable Images	6-6
6.3.2.1	Changing the Transfer Vector • 6-6	
6.3.2.2	Updating the Linker Options File • 6-6	
6.3.3	Updating Shareable Image Libraries	6-7

APPENDIX A VMS MODULAR PROGRAMMING STANDARD **A-1**

A.1	PURPOSE OF THIS STANDARD	A-1
A.2	APPLICABILITY	A-1
A.3	CODING RULES	A-2
A.3.1	The Calling Interface	A-2
A.3.2	Initialization	A-4
A.3.3	Reporting Exception Conditions	A-4
A.3.4	AST Reentrancy	A-5
A.3.5	Resource Allocation	A-5
A.3.6	The Format and Content of Coded Modules	A-6
A.3.7	Shareable Images	A-6
A.3.8	Upward Compatibility	A-7

APPENDIX B ARGUMENT CHARACTERISTICS **B-1**

B.1	VMS USAGE	B-1
B.2	DATA TYPE	B-6
B.3	ACCESS MECHANISM	B-8
B.4	PASSING MECHANISMS	B-8

INDEX

EXAMPLES

2-1	FORTRAN Program Showing the Improper Use of Implicit Arguments	2-6
2-2	FORTRAN Program Combining Procedures to Avoid Implicit Arguments	2-7
2-3	Static Storage and AST Reentrancy	2-14
2-4	A Sample Module Description	2-20
2-5	A Sample Procedure Description	2-22

Contents

3-1	A Sample Cross-Reference Listing Showing the References to the Symbol SPEED_OF_LIGHT	3-9
3-2	PASCAL Program Showing Use of Uppercase and Lowercase Characters in Code	3-10
3-3	BASIC Program Showing Use of Optional Spaces in Code	3-10
3-4	FORTRAN Program Showing Use of Block Comments in Code	3-10
3-5	PASCAL Program Which Uses a First-Time Flag	3-16
3-6	BASIC Initialization Procedure for LIB\$INITIALIZE	3-18
3-7	Program to Add Address to PSECT LIB\$INITIALIZE	3-18
3-8	BASIC Main Program	3-18
3-9	MACRO Program Showing Use of Queue Instructions to Perform All Accesses in a Single Instruction	3-22
3-10	MACRO Program Showing Use of Test and Set Instructions	3-23
3-11	A FORTRAN Program Disabling and Restoring ASTs	3-24
5-1	Transfer Vector Template	5-6
5-2	Template for a Linker Options File	5-8

FIGURES

1-1	The Software Life Cycle	1-2
1-2	Developing a Program that Calls Library Procedures	1-5
1-3	Procedures Available in the Run-Time Library	1-7
2-1	Levels of Abstraction	2-2
2-2	Possible Procedure Groupings	2-4
2-3	Designating Storage Responsibility to the Caller	2-9
2-4	Use of Storage Types	2-13
3-1	Examples of Facility Prefixes as Used in Procedure Names	3-2
3-2	Methods of Initializing	3-13
3-3	How to Initialize Static Storage	3-15
4-1	Black Box Testing Methods	4-3
4-2	White Box Testing Methods	4-4
4-3	A Sample Procedure for Integration Testing	4-5
5-1	Development of a User-Created Object Module Library	5-4
5-2	Creating a Shareable Image	5-5
6-1	Regression Testing	6-2
6-2	One Type of Argument Block, the Signal Argument Vector	6-5
B-1	Procedure Argument Passing Mechanisms	B-10

TABLES

2-1	Summary of Storage Use by Language _____	2-15
3-1	Common Library Facilities — Prefixes and Content _____	3-2
3-2	Naming Procedure Entry Points _____	3-4
3-3	Code for the Content and Usage of Global Variables _____	3-6
3-4	How to Declare Common Source Files _____	3-7
B-1	VMS Data Structures _____	B-1
B-2	VAX Standard Data Types _____	B-7
B-3	VAX Standard Passing Mechanisms _____	B-9

		(A) 1952
1-1	Amount of time in the day	1-1
1-2	Amount of time in the day	1-2
1-3	Amount of time in the day	1-3
1-4	Amount of time in the day	1-4
1-5	Amount of time in the day	1-5
1-6	Amount of time in the day	1-6
1-7	Amount of time in the day	1-7
1-8	Amount of time in the day	1-8
1-9	Amount of time in the day	1-9
1-10	Amount of time in the day	1-10

Preface

This manual is a tutorial guide to designing and coding modular procedures written in any VMS language. You can use these procedures for general programming. You can also include them in a procedure library, such as an object module library, shareable image, or shareable image library.

This manual includes required and optional modular programming techniques, recommended style, and a description of how to install modular procedures in procedure libraries.

A procedure is modular if it follows rules and principles that permit it to be successfully linked with other procedures that follow the same rules and principles. The *Guide to Creating VMS Modular Procedures* tells you how to follow the VMS Modular Programming Standard (which summarizes these rules and principles). Following this standard improves program reliability and reduces maintenance effort.

Intended Audience

This manual is intended for advanced system and applications programmers who are already familiar with VMS system concepts. Readers should be familiar with the VMS operating system and proficient in at least one supported language.

Document Structure

The general structure of this manual is based on the DIGITAL version of the software life cycle. All information is presented within this construct.

- Chapter 1 provides an overview of the software life cycle model used in this book. It also defines what is meant by the term "procedure," and provides an overview of procedures that already exist on the VMS operating system. The layered products that you might find useful in procedure development are also discussed.
- Chapter 2 provides information on design. The topics covered include organizing new applications, designing a modular procedure interface, using system resources, using input/output, writing internal documentation, and planning for signaling and condition handling.
- Chapter 3 presents general coding guidelines and information on initializing modular procedures. It also discusses guidelines for invoking optional user-supplied action routines, and writing AST-reentrant code.
- Chapter 4 describes methods used to test procedures for modularity, language-independence, and reentrancy. This chapter also provides general information on performance testing and monitoring procedures.
- Chapter 5 shows you how to create object module libraries, shareable images, and shareable image libraries from your completed procedures. Special attention is given to the transfer vector and the linker options file.
- Chapter 6 provides information about maintenance, such as upward compatibility, regression testing, updating procedures and procedure libraries, and changing the transfer vector or linker options file.

Preface

- Appendix A contains the VMS Modular Programming Standard.
- Appendix B presents the argument characteristics supported by the VAX Procedure Calling and Condition Handling Standard.

Associated Documents

The following documents are associated with this manual:

- *VMS Run-Time Library Routines Volume*
- *VMS System Services Reference Manual*
- *VMS Linker Utility Manual*
- All user's guides and reference manuals for all VMS languages

Conventions

Convention	Meaning
RET	In examples, a key name (usually abbreviated) shown within a box indicates that you press a key on the keyboard; in text, a key name is not enclosed in a box. In this example, the key is the RETURN key. (Note that the RETURN key is not usually shown in syntax statements or in all examples; however, assume that you must press the RETURN key after entering a command or responding to a prompt.)
CTRL/C	A key combination, shown in uppercase with a slash separating two key names, indicates that you hold down the first key while you press the second key. For example, the key combination CTRL/C indicates that you hold down the key labeled CTRL while you press the key labeled C. In examples, a key combination is enclosed in a box.
\$ SHOW TIME 05-JUN-1988 11:55:22	In examples, system output (what the system displays) is shown in black. User input (what you enter) is shown in red.
\$ TYPE MYFILE.DAT . . .	In examples, a vertical series of periods, or ellipsis, means either that not all the data that the system would display in response to a command is shown or that not all the data a user would enter is shown.
input-file, . . .	In examples, a horizontal ellipsis indicates that additional parameters, values, or other information can be entered, that preceding items can be repeated one or more times, or that optional arguments in a statement have been omitted.

Convention	Meaning
[logical-name]	Brackets indicate that the enclosed item is optional. (Brackets are not, however, optional in the syntax of a directory name in a file specification or in the syntax of a substring specification in an assignment statement.)
quotation marks apostrophes	The term quotation marks is used to refer to double quotation marks ("). The term apostrophe (') is used to refer to a single quotation mark.

1

Introduction to Modular Procedures

1.1 Following the Software Life Cycle

This manual discusses the development of modular procedures, procedure libraries, shareable images, and shareable image libraries, in terms of the software life cycle.

There are several different versions of the software life cycle that are currently in use. The one used in this manual depicts the software life cycle as having seven stages (see Figure 1-1). These stages are as follows:

- 1 The first stage is the concept stage. During the concept stage, the new application is proposed and considered.
- 2 The requirements or specifications stage is the point at which the goals of the new application are chosen and a requirements or specifications document is written outlining what will be accomplished by the new application.
- 3 At the design stage, detailed plans are made to ensure that the completed code will be modular and accurate, as well as to completely meet the goals decided upon at the requirements/specifications stage.
- 4 During the coding stage, the actual application is coded and internally documented.
- 5 Once coding is completed, the code is tested to ensure that it is accurate, efficient, and meets the requirements and specifications.
- 6 At the integration stage, the procedures and applications are grouped into facilities and libraries.
- 7 The maintenance stage is the point at which the completed applications are updated, enhanced, and modified.

Modularity becomes an important consideration in the design phase of the software life cycle. Because this manual is concerned primarily with modularity, detailed discussion begins at the design stage, and continues through the coding, testing, integration, and maintenance stages.

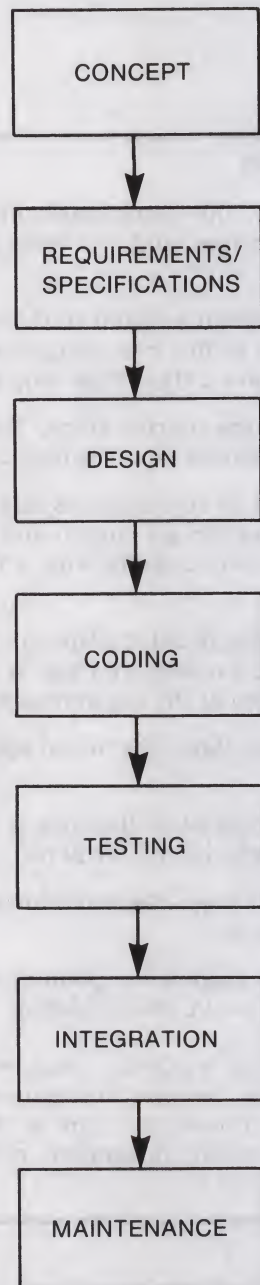
1.2 Defining the Modular Procedure

A procedure is a set of related instructions that performs a task. Typically, a procedure is invoked by executing a VAX CALLS or CALLG instruction. If you are using a high-level language, the compiler generates the CALLS or CALLG instruction on your behalf when you use the conventions required by your language to implement a procedure.

Introduction to Modular Procedures

1.2 Defining the Modular Procedure

Figure 1-1 The Software Life Cycle



ZK-3081-84

Introduction to Modular Procedures

1.2 Defining the Modular Procedure

VAX languages implement procedures as follows:

Ada	A procedure is declared as a PROCEDURE or FUNCTION.
BASIC	A procedure is a main program, subprogram, or FUNCTION.
BLISS	A procedure is declared as a ROUTINE.
C	A function is declared as an identifier followed by a balanced pair of parentheses which contain parameter declarations. There is no keyword PROCEDURE or FUNCTION in C. A procedure is declared as a function with a return type of "void." The return type of the function is specified before the identifier giving the name of the function.
COBOL	A procedure is a paragraph (or SECTION) or group of logically successive paragraphs (or SECTIONS) in the Procedure Division.
DIBOL	A procedure is a main program, subrouting, or function.
FORTRAN	A procedure is a main program, subroutine, or function.
MACRO	A procedure begins with an .ENTRY directive and ends with a RET instruction.
PASCAL	A procedure is declared as a PROCEDURE or FUNCTION.
PL/I	A procedure is an external or internal PL/I PROCEDURE or entry point.
RPG II	A procedure is a main program.
SCAN	A procedure is declared as a PROCEDURE or FUNCTION.

A procedure is modular if it follows rules and principles that permit it to be successfully linked together with other procedures that follow the same rules and principles. These rules and principles are summarized in the VMS Modular Programming Standard, contained in Appendix A.

This manual describes how to perform a complex task by dividing it into modules and coding each module as a separate procedure. This kind of modular programming offers several advantages over writing a complex program as a single source module.

- You can use any modular procedure in any program.
- You can add a modular procedure to a library at any time.
- You need not rewrite common algorithms every time they are needed for a new program.
- You can divide a complex program into simpler procedures to reduce development time and complexity, and increase reliability.
- You can modify or replace a procedure without modifying the calling program provided that you adhere to the optional guidelines for maintaining upward compatibility in the VMS Modular Programming Standard.
- You can control process-wide resource allocation.
- You can use different programming languages to write different procedures for a program.

Introduction to Modular Procedures

1.2 Defining the Modular Procedure

The VMS Modular Programming Standard also contains guidelines that are recommended rather than required. If you follow these recommendations, you will gain the following advantages in addition to modularity:

- Shareable library procedures can save memory space, disk space, and link time.
- AST-reentrant procedures can be called by AST-level procedures.
- Modular procedures that conform to all coding recommendations are similar in format and therefore are easier to use and maintain.

You can use modular procedures for general programming or you can group them in procedure libraries. Grouping procedures into libraries is simply a way of collecting procedures so that calling programs can access them easily. When you link your program to a library, the VMS linker automatically searches that library to resolve any references that your program makes to procedures in the library. Because the VMS Linker searches the specified library automatically, your program can call many modular procedures without having to include the name of each procedure explicitly in the LINK command. The program's executable image and the procedures that it calls are executed in the proper sequence at run time.

Figure 1-2 shows the development of a program that calls one or more procedures in a library. Depending on the options you select when writing modular procedures, you can control the way the linker accesses your procedures and thus the way procedures are invoked at run time. For example, if you place commonly used procedures within a shareable procedure library or shareable image library, you can save memory and disk space because all user processes may access a single copy of the shared procedures.

1.3 Existing VMS System Procedures

Many system routines that perform various advanced applications are included in the VMS operating system. Before you write a new procedure, you should check to make sure that the application does not already exist.

You may also find that the system procedures provide useful building blocks for your own procedures. Procedures designed to accomplish many general functions already exist in the system libraries; thus, your procedure can simply call an existing procedure rather than duplicate the code.

The following is a list of the four types of callable system procedures along with the titles of the manuals documenting those procedures:

Run-Time Library Procedures
VMS Run-Time Library Routines Volume

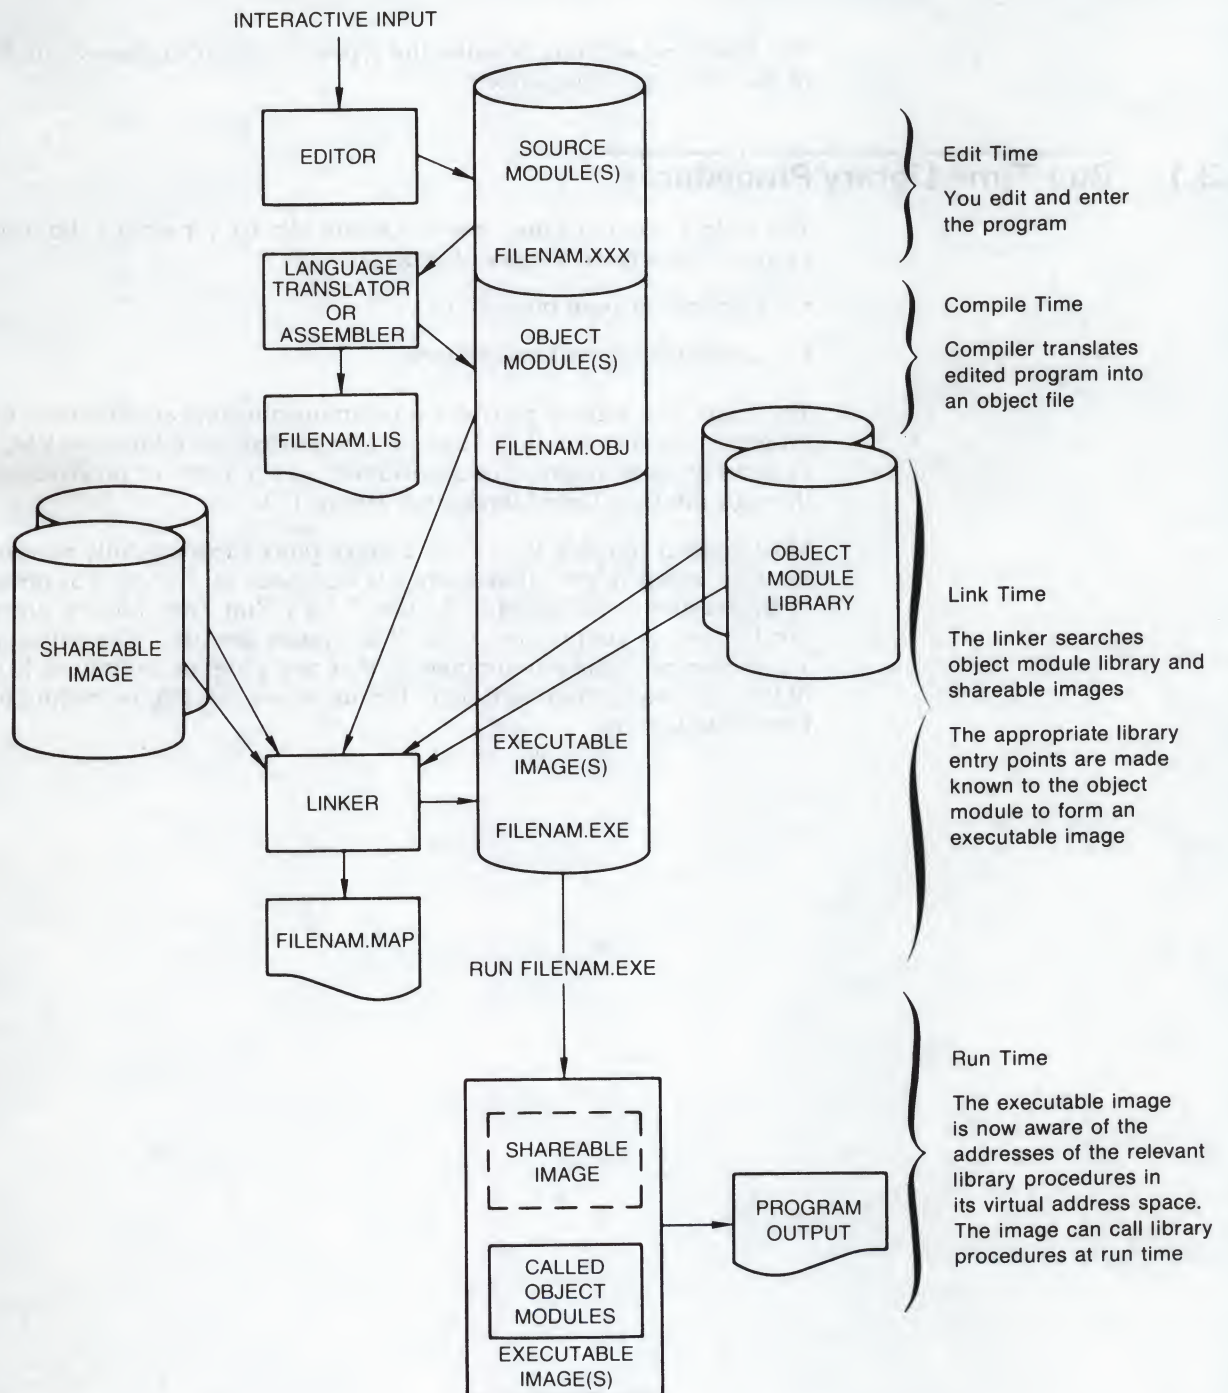
System Services
VMS System Services Reference Manual

Utility Routines
VMS Utility Routines Manual

Introduction to Modular Procedures

1.3 Existing VMS System Procedures

Figure 1-2 Developing a Program that Calls Library Procedures



ZK-4068-85

Introduction to Modular Procedures

1.3 Existing VMS System Procedures

Record Management Services
VMS Record Management Services Manual

The following sections describe the types of system routines available as part of the VMS operating system.

1.3.1 Run-Time Library Procedures

The VMS Common Run-Time Procedure Library (or simply the Run-Time Library) contains two types of procedures:

- General purpose procedures
- Language support procedures

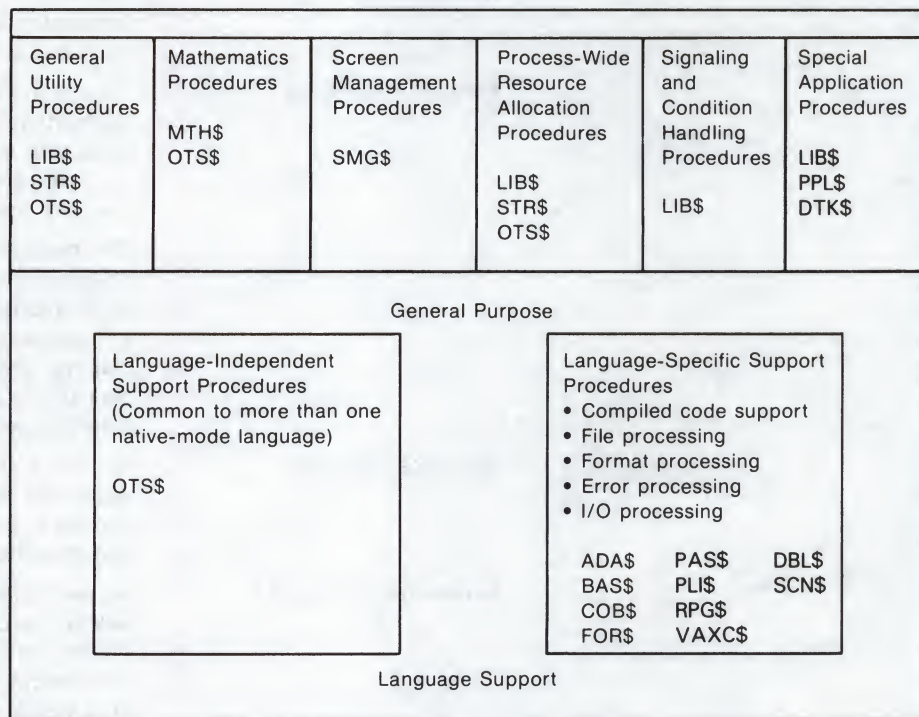
The Run-Time Library provides a common run-time environment for user programs because the Run-Time Library procedures follow the VMS Modular Programming Standard. For an overview of the types of procedures available through the Run-Time Library, see Figure 1-3.

Most general-purpose Run-Time Library procedures are fully reentrant and position independent. (Reentrancy is discussed in Section 3.3; position independence is discussed in Section 3.1.1.) Run-Time Library procedures can be used in conjunction with VMS system services. The advantage of a common run-time environment is that any program written in MACRO, BLISS, or a supported high-level language can call any procedure in the Run-Time Library.

Introduction to Modular Procedures

1.3 Existing VMS System Procedures

Figure 1-3 Procedures Available in the Run-Time Library



ZK-3082-84

The following chart describes the Run-Time Library procedures:

RTL Procedures	Function
DECtalk [™] management	DECtalk management procedures control DIGITAL's DECtalk device. Procedures are available to set device characteristics, speak text, read telephone keypad input, and other auxillary functions.
General utility	This group includes procedures for getting a record from a device, performing string manipulation, converting data types for input and output, and obtaining the system date or time.

[™]DECtalk is a trademark of Digital Equipment Corporation.

Introduction to Modular Procedures

1.3 Existing VMS System Procedures

RTL Procedures	Function
Mathematics	Mathematics procedures perform common arithmetic, algebraic, and trigonometric functions, such as calculating the sine of an angle.
Parallel processing	<p>This group of routines provides support for an application that can run as several threads of execution in parallel. Procedures are available to create threads, schedule work, and share memory, mutexes, and semaphores.</p> <p>The parallel processing procedures perform operations typically required in a parallel application, such as subprocess creation, interprocess communication, and resource sharing. Although designed for parallel processing applications, these procedures can also be used on uniprocessor systems.</p>
Resource allocation	Resource allocation procedures allocate and deallocate virtual memory, VMS local event flag numbers, BASIC/FORTRAN logical unit numbers, and dynamic strings.
Screen management	<p>Screen management procedures perform terminal-independent screen management functions.</p> <p>These procedures assist you in designing, composing, and keeping track of complex images on a video screen.</p>
Signaling and condition handling	These procedures perform operations involved with handling exception conditions such as signaling exceptions, establishing condition handlers, and enabling the detection of hardware exceptions.
Language-independent support	Language-independent support procedures can be called from any VAX language. Although most of these routines are in the OTS\$ facility, some LIB\$ routines also perform language-support operations. Language-independent support procedures are mostly data-type conversion procedures.
Special applications	Procedures for specialized applications, such as syntax analysis and cross-reference tasks, are also available in the Run-Time Library.

1.3.2 System Services

System services are procedures that the VMS operating system uses to control resources available to processes. They are also used to provide for communication among processes and to perform basic operating system functions such as the coordination of input/output operations.

Although most system services are used primarily by the operating system itself on behalf of logged-in users, many are available for use in application programs. For example, when you log in to the system, the Create Process system service (\$CREPRC) is called to create a process on your behalf. You may, in turn, write a program that calls \$CREPRC to create a subprocess.

Introduction to Modular Procedures

1.3 Existing VMS System Procedures

System services can be divided into functional groups. The following list describes the function of each group of system services:

Service Group	Function
Security	The security services provide various mechanisms that you can use to enhance the security of VMS systems.
Event flag	Event flags can be used by a process to synchronize sequences of operations in a program. Event flag services clear, set, and read event flags, and place a process in a wait state pending the setting of an event flag or flags.
AST	Process execution can be interrupted by events (such as I/O completion) to allow for the execution of designated subroutines. These software interrupts are called asynchronous system traps (ASTs) because they occur asynchronously to process execution. There are system services that allow a process to control the handling of ASTs.
Logical names	Logical name services are used to maintain and access character-string logical name and equivalence name pairs. Logical names can provide device independence for system and application program input and output operations.
Input/output	<p>I/O services perform input and output operations directly, rather than through the file handling services of the VMS Record Management Services (RMS). I/O services do the following:</p> <ul style="list-style-type: none">• Perform logical, physical, and virtual input/output operations• Perform network operations• Queue messages to system processes
Process control	Process control services allow you to create, delete, and control the execution of processes.
Timer and time conversion	Timer services schedule program events for a particular time of the day, or after a specified interval of time has elapsed. The time conversion services provide a way to obtain and format binary time values for use with the timer services.
Condition handling	Condition handlers are procedures that receive control when a hardware or software exception condition occurs during image execution. Condition-handling services designate condition handlers for special purposes.

Introduction to Modular Procedures

1.3 Existing VMS System Procedures

Service Group	Function
Memory management	Memory management services provide ways to use the virtual address space available to a program. Included are services that do the following: <ul style="list-style-type: none">• Allow an image to increase or decrease the amount of virtual memory available• Control the paging and swapping of virtual memory• Create and access files in memory that contain shareable code or data
Change mode	Change mode services alter the access mode of a process. These services are used primarily by the operating system.
Lock management	Lock management services allow cooperating processes to synchronize their access to shared resources.

1.3.3 Utility Routines

Utility routines perform a particular task or set of tasks. For example, the Print Symbiont Modification (PSM) routines can be used to modify the VMS print symbiont, and the EDT routines can be used to invoke the EDT editor from a program.

When using a set of utility routines that performs the same tasks as a VMS utility, you should read the documentation for that utility. Doing so will provide you with additional information on the tasks that each set of routines can perform. The following list shows VMS utilities that have corresponding utility routines:

Routine Group	Function
Command Language (CLI)	CLI routines process command strings using information from a command table. A command table is a list of definitions that describe the allowable formats for commands.
Convert and Convert/Reclaim (CONV)	CONV routines copy records from one or more files to an output file, changing the record format and file organization to that of the output file. The Convert /Reclaim routine reclaims empty buckets in Prolog 3 indexed files so that new records can be written into them. These routines cannot be called from AST level.
File Definition Language (FDL)	FDL routines perform many of the functions of the RMS File Definition Language. These routines cannot be called from AST level.

Introduction to Modular Procedures

1.3 Existing VMS System Procedures

Routine Group	Function
Librarian (LBR)	LBR routines are used to create and maintain libraries and library modules. Libraries are files that provide a convenient way to organize frequently used modules of code or text.
National Character Set (NCS)	NCS routines provide a common facility for defining collating sequences and conversion functions.
Sort/Merge (SOR)	SOR routines are used to implement a sort or merge operation within a program. These routines are reentrant; that is, a number of sort or merge operations can be active at the same time.
Terminal Fallback Facility (TFF)	TFF routines provide table-driven character conversion for terminals.

The following facilities are accessible only through the utility routines; there are no corresponding VMS utilities:

Routine Group	Function
Data Compression/Expansion (DCX)	DCX routines are used to analyze and compress data records and then expand the compressed records to their original state. No information is lost in this compression /expansion process.
Editor (EDT)	The EDT routine invokes the EDT editor from within a program. The program may be written to handle the editing work, or EDT may run interactively to allow a user at the terminal to edit a file while the program is running.
Print Symbiont (PSM)	PSM routines modify the behavior of the print symbiont that is supplied with the VMS operating system.
Symbiont/Job Controller Interface (SMB)	SMB routines provide the interface between the job controller and symbiont processes.

1.3.4 Record Management Services

The VMS Record Management Services (RMS) assist user programs in processing and managing files and their contents. RMS allows you to create files that use a minimum amount of system resources while decreasing input/output time.

Introduction to Modular Procedures

1.3 Existing VMS System Procedures

The two types of Record Management Services are as follows:

RMS Service Group	Function
File-related services	File-related services create and access new files, access existing files, extend the disk space allocated to a file, close a file, obtain file characteristics, and perform other functions related to the file as a whole.
Record-related services	Record-related services get, locate, insert, delete, and update records, as well as other record-related operations.

For complete descriptions of these Record Management Services, see the *VMS Record Management Services Manual*.

1.4 Tools for Developing Applications on VMS

There are several tools specifically designed to aid you in developing applications in the VMS environment.

- VAX DEC/CMS

The Code Management System (CMS) is a program library system for software development and maintenance. CMS works as an online librarian for a project. While project members perform the normal functions of program development, modification and testing, CMS keeps a record of changes made to their files.

- VAX DEC/MMS

The Module Management System (MMS) is a tool that automates and simplifies the building of software systems. MMS is useful for building both simple programs, which may have only one or two source files, and complex programs, which may consist of several source files, message files, and documentation. It can rebuild all the components in a system, or rebuild only those that have changed since the system was last built.

- VAX DEC/Test Manager

The VAX DEC/Test Manager is a tool that organizes software tests and automates the way you run tests and evaluate the results. It provides an efficient way to organize, run, and store the results of exiting tests. The VAX DEC/Test Manager is based on the concept of regression testing.

- VAX Language-Sensitive Editor

The VAX Language-Sensitive Editor is a source code editor that allows you to quickly and accurately develop programs using the EDT commands. The VAX Language-Sensitive Editor has useful development features such as templates, placeholders, user-defined templates and placeholders, and commands. Other features include a compiler interface and online language help.

- VAX Performance and Coverage Analyzer

The VAX Performance and Coverage Analyzer is a tool that can be used by software developers to gather performance or test coverage data on user programs. This tool will then present that data in tables, bar histograms, annotated source listings, and other formats.

Introduction to Modular Procedures

1.4 Tools for Developing Applications on VMS

- VAX Source Code Analyzer (SCA)

The VAX Source Code Analyzer (SCA) is an interactive, multilanguage, source code cross-reference and static analysis tool designed to aid developers in understanding the complexities of large-scale software systems. The VAX Language-Sensitive Editor and SCA together provide an integrated method for creating, compiling, correcting, and inspecting source code within a single session.

These tools are optional products. They are not included with VMS. For further information on these products, see your DIGITAL sales representative.

REPORT OF THE COMMISSIONER OF THE
BUREAU OF LAND MANAGEMENT

TO THE SECRETARY OF THE INTERIOR

FOR THE YEAR 1904

WASHINGTON, D. C.

2 Design

Well-designed procedures are more likely to be modular, well written, and easy to maintain. Any time that is saved by skimping at the design stage is lost several times over in patchwork attempts to remedy problems stemming from a poor design.

This chapter discusses the following aspects of designing a new application:

- Organizing new applications
- Designing a modular procedure interface
- Using system resources
- Using I/O
- Beginning the internal documentation
- Planning for signaling and condition handling

2.1 Organizing New Applications

The first work to be done in designing a new application is to look at the overall organization. An application should be made up of one or more files, each containing one or more procedures. When linked, the procedures are organized into program sections (PSECTs). Each procedure, as well as the interface between the procedures, must be designed to conform to the VMS Modular Programming Standard (see Appendix A).

2.1.1 Organizing Files and Modules

Each application contains one or more files. Each file contains exactly one module. For information on naming files, refer to Section 3.1.2.3. For information on naming modules, refer to Section 3.1.2.4.

2.1.2 Organizing Procedures into Modules

Each module should contain a single procedure or a group of related procedures. The VMS Linker always brings the entire module containing a called procedure into the image if any of its entry points are referenced. Thus, placing each procedure in a separate module reduces image size. It also allows more flexibility when using a procedure library because you can supply your own version of one procedure while using other procedures from the library. If many procedures have been grouped in a single module, the linker must link all or none of them.

You should group procedures into a module if they share the same static storage or if they have a similar calling sequence, perform similar functions, and share a significant amount of code.

Design

2.1 Organizing New Applications

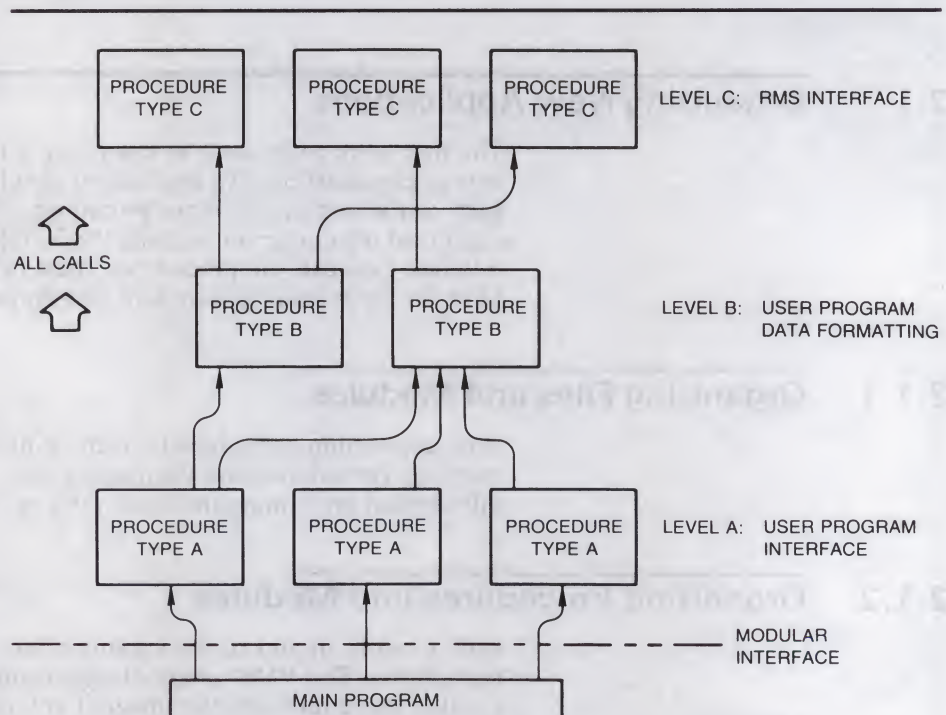
If you are writing a large number of related procedures that call one another or access common data blocks, you should make the relationship among those procedures as clear as possible. You can do this by organizing those procedures to minimize the interaction between procedures, and between procedures and data structures. There are several guidelines to help you minimize such interaction:

- Organize procedures into levels of abstraction.
- Make sure each level calls only the next lower level.
- Restrict read/write access to data structures and system components to as few procedures as possible.

Figure 2-1 shows the BASIC and FORTRAN record I/O processing procedures. These are implemented in the following three levels of abstraction:

- 1 User program interface (UPI)
- 2 User program data formatting (UDF)
- 3 Record processing and VMS RMS interface (REC)

Figure 2-1 Levels of Abstraction



ZK-4006-85

2.1 Organizing New Applications

All calls are made in one direction, to the next innermost level. Procedures at different levels should be in different modules. Figure 2-2 shows possible groupings of procedures.

2.2 Defining a Modular Procedure Interface

Procedures communicate with one another by passing arguments. In order to clarify the interactions between procedures and programs, each argument must be defined in the design stage.

2.2.1 Explicit Arguments

Explicit arguments are a procedure's primary interface with other programs. Therefore, rules for argument order, data types, and passing mechanisms must be followed carefully to maintain a modular interface. The following format is used to describe each argument:

argument-name

VMS usage:	argument-data-structure
type:	argument-data-type
access:	argument-access
mechanism:	argument-passing-mechanism

(See Appendix B for descriptions of each of these four argument attributes.)

To make your procedures easier to call, the passing mechanism used for particular data types should be consistent throughout all procedures in a facility. Passing all atomic data by reference and all string data by descriptor is recommended.

2.2.2 Implicit Arguments

An implicit argument is one that is not specified in the argument list. Implicit arguments provide additional information to your procedure from static storage locations. There are two types of implicit arguments:

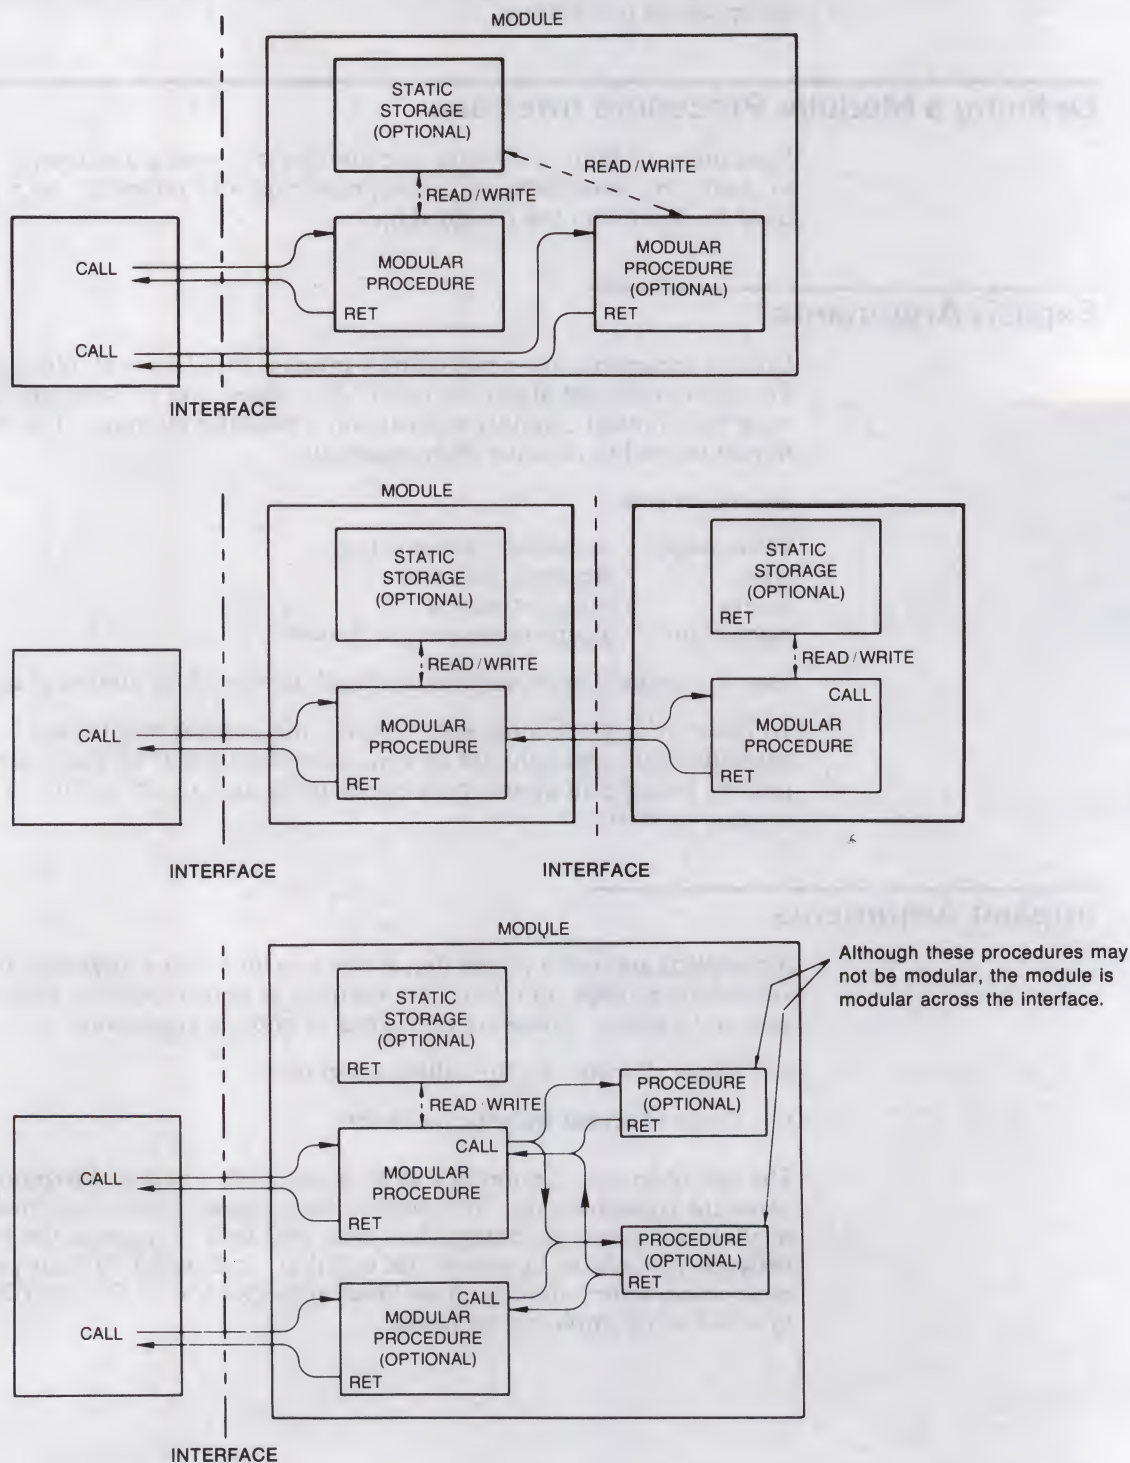
- Those allocated by the calling program
- Those allocated by your procedure

The use of implicit arguments violates the VMS Modular Programming Standard (Appendix A). Their use is discouraged because they make the relationship across procedures less clear and tend to increase the interaction between procedures in a way that might go undetected. If your procedure must retain information from previous activations, read Section 2.2.3 for ways to avoid using implicit arguments.

Design

2.2 Defining a Modular Procedure Interface

Figure 2-2 Possible Procedure Groupings



ZK-4007-85

2.2 Defining a Modular Procedure Interface

2.2.2.1 Implicit Arguments Allocated by the Calling Program

The calling program can allocate implicit arguments as statically allocated variables in a named PSECT (for example, COMMON and MAP in BASIC, COMMON in FORTRAN, or variables declared in the outer block of a procedure or program in PASCAL). The calling program can also allocate implicit arguments as statically allocated global variables (for example, symbols defined with a double colon `::` in MACRO and GLOBAL variables in BLISS).

Allocation of implicit arguments by the calling program violates the VMS Modular Programming Standard for the following reasons:

- Two programs could use the same PSECT name or global variable for different values. This error would be undetected.
- The calling program is no longer independent of the called procedure. Consequently, a change in one could inadvertently affect the other.
- In FORTRAN, the calling program has to declare all variables as COMMON regardless of the number of implicit inputs actually needed. All COMMON variables should also be declared by all modules that use the COMMON storage, further decreasing independence.

2.2.2.2 Implicit Arguments Allocated by the Called Procedure

Implicit arguments allocated by the called procedure are kept in local static storage.

These implicit arguments are usually used to keep track of resources (using resource allocating procedures) and shorten the explicit argument list. However, the use of implicit inputs by non-resource-allocating procedures can lead to unexpected results. For example, assume that procedure A is to leave information for a companion procedure B. This would result in B having both explicit inputs (from its caller) and implicit inputs (from A's storage). Next, consider that a calling program calls A, then calls procedure X, and finally calls B. For the calling program to get correct results from B, it must know that X (and any procedure that X calls) did not make a call to A, since such a call would change the implicit inputs A leaves for B.

Since one of the objectives of modular programming is to permit procedures to be combined arbitrarily without needing to understand each other's internal workings, the use of such implicit arguments violates the VMS Modular Programming Standard. The same problems can occur with any non-resource-allocating procedure that leaves results for itself as future implicit arguments.

2.2.3 How to Avoid Using Implicit Arguments

Procedures that do not allocate resources can be written in the following three ways to avoid the implicit argument problems described in Section 2.2.2:

- When one procedure obtains results from another, combine the two procedures into a single call. (See Section 2.2.3.1.)
- Provide a single call to an action routine that is supplied by the calling program part way through the procedure's execution. (See Section 2.2.3.2.)

Design

2.2 Defining a Modular Procedure Interface

- Give the calling program responsibility for retaining information from a procedure activation. This is done with an explicit argument. (See Section 2.2.3.3.)

2.2.3.1 Combining Procedures

Often, non-resource-allocating procedures can be combined into a single procedure that returns all information explicitly in a single call.

Compare Example 2-1 with Example 2-2. to see the effects of combining procedures to avoid the use of implicit arguments.

Example 2-1 FORTRAN Program Showing the Improper Use of Implicit Arguments

```
!+
! This program demonstrates a situation where
! the input of a procedure depends on the output
! of a previously called procedure.
!-
      REAL*4 X, Y, RESULT
      X = 1
      Y = 1

!+
! Call the procedure that writes into a common data area.
!-
      CALL SUM_SQUARES (X, Y)

!+
! Call the procedure that reads from the common data area.
!-
      CALL GET_SQRT (RESULT)

!+
! Print the result obtained.
!-
      WRITE (6,10) X, Y, RESULT
10    FORMAT(1X, 'SQRT(', F6.2, '**2 + ', F6.2, '**2) =', F6.2)
      STOP
      END

!+
! This procedure sums the squares of its two inputs and
! places the result in a common area, for use by some
! other procedure.
!-
      SUBROUTINE SUM_SQUARES (A, B)
      COMMON /INTERNAL_STORAGE/ TEMP_RESULT
      TEMP_RESULT = (A ** 2) + (B ** 2)
      RETURN
      END

!+
! This procedure calculates the square root of whatever
! number is in the common area.
!-
      SUBROUTINE GET_SQRT (C)
      COMMON /INTERNAL_STORAGE/ TEMP_RESULT
      C = SQRT (TEMP_RESULT)
      RETURN
      END
```

2.2 Defining a Modular Procedure Interface

Example 2-2 FORTRAN Program Combining Procedures to Avoid Implicit Arguments

```

!+
! This procedure shows the subroutines called in
! the previous example combined into a single subroutine
! that eliminates the use of COMMON.
!-
      REAL*4 X, Y, RESULT
      X = 1
      Y = 1
!+
! Call the new procedure.
!-
      CALL DO_IT_ALL (X, Y, RESULT)
      WRITE (6,10) X, Y, RESULT
10    FORMAT (1X, 'SQRT (', F6.2, '**2 + ', F6.2, '**2) = ', F6.2)

      STOP
      END

!+
! This procedure calculates the square root of the sum of
! the squares of its first two arguments, and returns the
! result in the third argument. It combines the functions
! provided by the SUM_SQUARES and GET_SQRT
! procedures and eliminates the use of COMMON.
!-
      SUBROUTINE DO_IT_ALL (A, B, C)
      C = SQRT ((A ** 2) + (B ** 2))
      RETURN
      END

```

2.2.3.2 User Action Routine

Another way to combine several procedures into one call is to let the calling program gain control at a critical point in your procedure's execution. For this to happen, your procedure must specify an action routine argument that is called during execution. Thus, your procedure can execute twice, before and after the action routine, with no implicit inputs. The OPEN statements in BASIC, FORTRAN, and PASCAL use this technique by permitting the user to supply a user action routine.

To keep the calling program from having to provide implicit inputs for its action routine, your procedure should also provide another argument that is passed to the action routine. The calling program uses the following calling sequence to invoke your procedure:

```
CALL my-proc (... ,action-routine ,user-arg)
```

Then your procedure invokes the action routine as follows:

```
CALL action-routine (... ,user-arg)
```

For information on writing user-action routines, see Section 3.1.6.

Design

2.2 Defining a Modular Procedure Interface

2.2.3.3 Designating Responsibility to the Calling Program

You can make the calling program responsible for retaining information from one procedure activation to another. There are three ways to accomplish this:

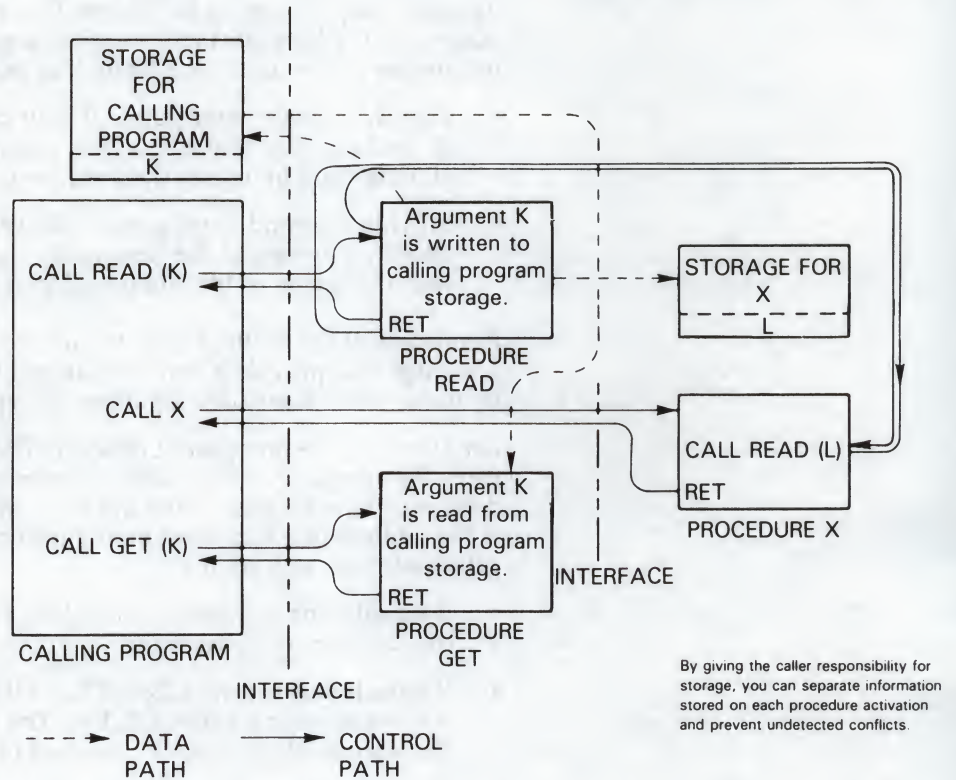
- Require the calling program to allocate the storage your procedure needs. Then have the calling program pass the address of the storage location as an explicit argument on all calls to your procedure. The disadvantage of this method is that you cannot increase the amount of storage needed by your procedure without requiring all calling programs to be rewritten. Thus, you should use this method only when you are confident that your procedure will not be revised to use additional storage in the future.
- Require the calling program to allocate a longword pointer to the stored data and pass its address to your procedure as an explicit argument. On the first call, your called procedure will dynamically allocate storage (by calling `LIB$GET_VM`) and store its address in the caller's longword. On subsequent calls, your procedure will use information left in the storage area from previous calls.
- Require the calling program to pass a processwide identifying value to your procedure on all calls. The processwide identifier indicates which information from previous procedure activations is to be used as implicit inputs.

Figure 2-3 shows a calling program that has responsibility for explicitly indicating the storage to be used by the called procedure.

Design

2.2 Defining a Modular Procedure Interface

Figure 2-3 Designating Storage Responsibility to the Caller



ZK-4004-85

Calling Program Allocates Procedure Storage

This method causes the calling program to allocate all storage needed and pass the address of the storage as an explicit argument on each call.

For example, the library procedure **MTH\$RANDOM** requires that the calling program allocate storage for the longword seed and pass its address on each call. **MTH\$RANDOM** takes the seed as input and computes the next random number sequence from the current seed value. **MTH\$RANDOM** returns a random number between 0 and 1 and updates the longword seed passed by the calling program. This ensures that the procedure will generate a different value on the next call.

The next two sections describe interface techniques that permit storage size to change without affecting the interface with the calling program.

Design

2.2 Defining a Modular Procedure Interface

Calling Program Passes Pointer

In this method, the calling program allocates only a longword pointer to the dynamic heap storage to be allocated by your procedure. It then passes the address of the longword as an explicit argument. The following two interface techniques can be used to indicate that storage is to be initialized:

- Provide a single entry point. If your called procedure finds the value zero in the longword that the calling program has allocated, the procedure allocates and initializes dynamic heap storage.
- Provide a second entry point. This entry point stores the address of the allocated storage in the longword. On subsequent calls, your procedure uses that value as the storage address of information from previous calls.

Regardless of the method used to indicate storage allocation and initialization, you must also provide a way to indicate storage deallocation. You can do this by using either a separate argument or separate entry point.

For example, the procedure `LIB$INIT_TIMER`, which gets times and counts from the operating system, uses a single optional argument **handle-adr** to determine where these values are to be stored. The **handle-adr** argument is the address of a longword pointing to a block of storage that contains the values of times and counts.

- If **handle-adr** is missing, the values are stored in static storage, making this call non-AST-reentrant.
- If **handle-adr** is zero, `LIB$INIT_TIMER` allocates a block of dynamic heap storage by calling `LIB$GET_VM`. The values are placed in that block, and the address of the block is returned in **handle-adr**.
- If **handle-adr** is nonzero, it is considered to be the address of a storage block previously allocated by a call to `LIB$INIT_TIMER`. The block is then used again and new times and counts are stored in it.

`LIB$FREE_TIMER` deallocates the block of dynamic heap storage allocated by a previous call to `LIB$INIT_TIMER`. The **handle-adr** argument to `LIB$FREE_TIMER` is the address of a longword that points to a block of dynamic heap storage where times and counts have been stored. That storage is returned to free storage by calling `LIB$FREE_VM`.

Calling Program Passes a Processwide Identifier

In this method, the calling program passes a processwide identifying value to identify implicit results produced on previous calls, which will be implicit inputs on this call. Any calling program can use the processwide identifier. Examples include BASIC or FORTRAN logical unit numbers and VMS system services I/O channel numbers.

Processwide identifiers are a resource. Modular programming techniques require that all resources allocated by a procedure be allocated by calling a resource-allocating procedure. This prevents conflicts because a single procedure can keep track of multiple allocations to more than one procedure or procedure activation. Therefore, if you use the method described in this section, you will also have to write a resource-allocating procedure to control the resource. If you write a resource-allocating procedure, it is recommended that you place it in an object module library so that other programmers can use it.

2.2 Defining a Modular Procedure Interface

The library procedures LIB\$GET_LUN and LIB\$FREE_LUN allocate and deallocate FORTRAN and BASIC logical unit numbers outside the range normally specified in user programs, that is, outside the range 0 to 99.

2.2.4 Order of Arguments

Procedures in the Run-Time Library follow a consistent pattern for positioning arguments. You should follow the same guidelines. Group procedure arguments from left to right in the following order:

- 1 Required input arguments (read access)
- 2 Required input-output arguments (modify access)
- 3 Required output arguments (write access)
- 4 Optional input arguments (read access)
- 5 Optional input-output arguments (modify access)
- 6 Optional output arguments (write access)

Note that optional arguments follow required arguments. Therefore, when the calling program omits the optional arguments, the actual argument list passed to the procedure is shortened.

The called procedure accesses the required arguments from left to right, beginning with the first argument. The only exceptions are procedures that return a function value longer than 64 bits, such as strings or H_floating values. Most function values are returned in the first two registers, R0/R1. (This is done automatically by high-level languages.) However, when the function value exceeds 64 bits, it is too large to be returned in R0/R1. In this case, the calling program uses the first argument to specify where the function value is to be stored, and the other arguments are shifted right one position. (See the VAX Procedure Calling and Condition Handling Standard in the *Introduction to VMS System Routines*.)

2.2.5 Using Optional Arguments

An optional argument is one that the calling program can omit. The calling program indicates the omission by passing argument list entries containing zero. For a trailing optional argument, the calling program can pass a shortened list or a zero argument list entry.

A zero argument list entry is simply a zero passed to the procedure by value. For example, if we call a procedure called GRA_CUBE and omit an optional argument C, the calling sequence from BASIC would be as follows:

```
15 CALL GRA_CUBE(A, B, 0 BY VALUE)
```

In this call, "0 BY VALUE" is the zero argument list entry.

Note: Most VMS system services, unlike the Run-Time Library procedures, cannot accept a shortened argument list. Omitted arguments must always be indicated with a zero argument list entry. For arguments passed by value, there is no distinction between passing a zero value and passing a zero argument list entry.

Design

2.3 JSB Entry Points

2.3 JSB Entry Points

DIGITAL recommends that you do not use JSB¹ entry points in procedures that will be contained in a procedure library. Procedures that can be invoked only by JSB instructions are not callable by high-level languages and violate the VAX Procedure Calling and Condition Handling Standard. If a procedure does use a JSB entry point, it must also provide an equivalent call entry point to maintain language independence. The call entry point must be provided because JSB instructions are only available in MACRO and BLISS32.

If you provide a JSB entry point for your procedure, the name of the JSB entry point is the same as the name of the procedure, except that it ends in `_Rn`. The *n* indicates the highest register modified or used as an input argument.

For example, the JSB entry point of the Run-Time Library procedure `LIB$ANALYZE_SDESC` is `LIB$ANALYZE_SDESC_R2`.

2.4 Using System Resources

The system resources available to you are limited by your account quotas and by the amount of available resources on the system. Efficient use of system resources makes more resources available for all processes.

2.4.1 Choosing a Storage Type

There are three types of storage: static, stack, and heap. The three forms of storage differ in the method and duration of allocation, that is, how long that storage is in use.

2.4.1.1 Stack Storage

A procedure dynamically allocates stack storage on the process stack at run time, as needed. To allocate stack storage, the procedure moves the stack pointer "up" by decreasing its value. Note that stack storage is not initialized to zero because the stack is created once and reused many times for subsequent stack frames.

The procedure deallocates stack storage by moving the stack pointer "down" (increasing its value) when that procedure returns control to the calling program. Stack storage exists only for the duration of the procedure activation that creates it.

2.4.1.2 Heap Storage

Dynamic heap storage is allocated at run time from a processwide pool, as the procedure activation needs it and as the account quotas and virtual address space of your process permits.

To allocate heap storage, your procedure calls a system routine such as the Run-Time Library procedure `LIB$GET_VM` or the system service `$EXPREG`. The call to the system routine may be within the procedure itself, or you may use a general resource-allocating procedure to centralize your resource allocations.

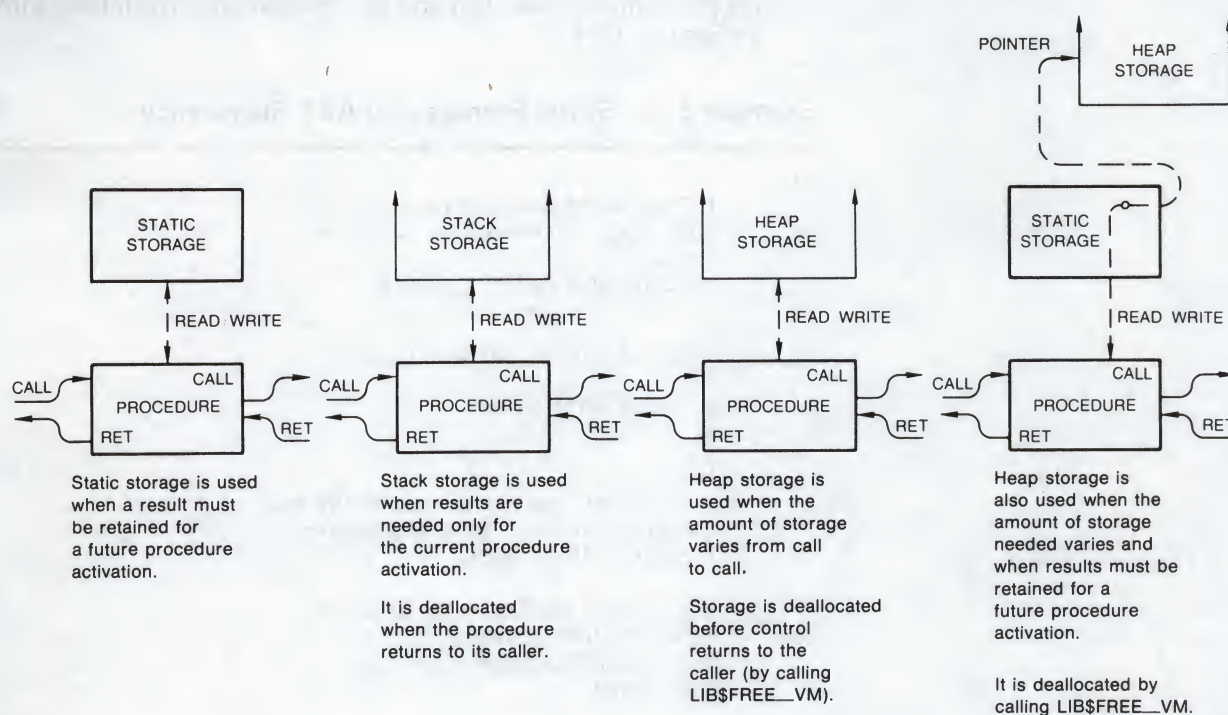
Heap storage is deallocated—that is, returned to the processwide pool—by calling `LIB$FREE_VM`. The system service `$CNTREG` cannot be used to deallocate heap storage.

¹ JSB is a MACRO instruction that means jump to subroutine.

Figure 2-4 shows how the different types of storage are used.

Note: The type of storage to be used can be determined by the duration or quantity of the storage. Any storage that is of long duration and unknown quantity (at compile time) should be heap storage. Storage of short duration (during the current invocation of the procedure) should be stack storage. Storage of long duration that is needed in only one instance should be static storage.

Figure 2-4 Use of Storage Types



ZK-4005-85

2.4.1.3 Static Storage

At link time, the VMS Linker collects storage in similar PSECTs into a single image section. The initial contents of this storage are specified in the source program. The VMS operating system initializes any non-initialized static storage to zero. On calls to a procedure after initialization, the static storage has the same allocation and the contents left from the previous call.

Design

2.4 Using System Resources

2.4.1.4 Avoiding Use of Static Storage

There are several disadvantages to using static storage:

- It is an inefficient use of memory. When using static storage, you must provide for the largest possible memory use.
- An image size is larger because of the inefficient use of memory.
- It can easily lead to problems with AST reentrancy, as seen in Example 2-3. This example circumvents the problem of an AST corrupting data by setting a first-time flag. Another method of preventing this problem is to use "test and set" instructions. For further information, see Section 3.3.4.2.

Example 2-3 Static Storage and AST Reentrancy

```
10      !+
      ! Program to demonstrate corruption
      ! of static storage due to AST's.
      !-
      DECLARE LONG CURRENT_NUMBER

      !+
      ! Enable CTRL/C AST handling.
      !-
      ON ERROR GOTO 19000
      X% = CTRLC

      !+
      ! Increment the number and print the
      ! current value. When the number
      ! reaches 1000, exit.
      !-
      FOR CURRENT_NUMBER = 1% TO 1000%
100     PRINT CURRENT_NUMBER;
      NEXT CURRENT_NUMBER
      GOTO 32767

19000   !+
      ! Error-handling routine. If this routine is
      ! entered due to a CTRL/C
      ! AST, corrupt CURRENT_NUMBER by setting it to -1.
      !-
      IF ERR = 28 THEN CURRENT_NUMBER = -1%
      RESUME 100

32767   END
```


2.4.1.5 Summary of Storage Use by Language

Table 2-1 summarizes storage available to the programmer in various language procedures.

Table 2-1 Summary of Storage Use by Language

Language	Storage Type		
	Static	Stack	Heap
Ada	Constants and fixed-size objects contained in library packages	Local subprogram and task variables	Dynamically sized objects in library packages and objects created by allocators
BASIC	All COMMON and MAP data storage Most arrays	Local variables Executable DIMENSION statement	Dynamic strings
BLISS	OWN and GLOBAL	STACK LOCAL	By calling LIB\$GET_VM
C	Objects declared with external or static internal linkage	Objects declared inside a function with "automatic" linkage	By calling malloc, calloc, or realloc
COBOL	All data storage	Not applicable	By calling LIB\$GET_VM
DIBOL	All RECORD, COMMON, and LITERAL data storage	Not applicable	Not applicable
FORTRAN	All data storage	Not applicable	By calling LIB\$GET_VM
MACRO	Block storage	Decrementing stack pointer	By calling LIB\$GET_VM
PASCAL	All program or module level storage	PROCEDURE and FUNCTION local	By calling NEW ¹
PL/I	STATIC	AUTOMATIC	ALLOCATE statement (BASED) ²
RPG II	All data storage	Not applicable	By calling LIB\$GET_VM
SCAN	STATIC, GLOBAL, COMMON, EXTERNAL	When AUTOMATIC is used in a procedure or macro	DYNAMIC STRING values, TREE pointers, and the ALLOCATE function

¹Although this is true most of the time, there are other rules that can also determine STATIC versus STACK allocation. For further information, see the VAX PASCAL User's Guide.

²BASED is the storage class used to allocate heap storage in PL/I. The ALLOCATE statement does the actual allocation.

Design

2.4 Using System Resources

2.4.2 Using Event Flags

Event flags allow modular procedures to communicate with each other and to synchronize their operations. Because they can be allocated at run time, event flags allow one procedure to run independently of other procedures existing in the same process.

Event flags are allocated and deallocated by the Run-Time Library procedures `LIB$GET_EF` and `LIB$FREE_EF`. (For further information, see the *VMS Run-Time Library Routines Volume* and the descriptions of the `LIB$GET_EF` and `LIB$FREE_EF` procedures which appear in Part II of that manual.)

2.4.3 Using Logical Unit Numbers

A logical unit number is used to define the device or file a program uses to perform input and output. Modular procedures do not need to know the unit numbers of other procedures running at the same time.

Logical unit numbers are used only in BASIC and FORTRAN.

Logical unit numbers should be allocated and deallocated using the `LIB$GET_LUN` and `LIB$FREE_LUN` Run-Time Library procedures. (For further information on using logical unit numbers, see the descriptions of the `LIB$GET_LUN` and `LIB$FREE_LUN` procedures in Part II of the *VMS Run-Time Library Reference Manual*.)

2.5 Using Input/Output

In general, your procedure's input/output (I/O) is directed to either the terminal or a file. (In some cases, you may need to use mailbox I/O and network operations. For information on these areas, see *VMS Networking Manual*.) Regardless of whether you are directing input/output to the terminal screen or to a file, there are two rules you must follow to maintain modularity:

- A procedure must not print error or informational messages either directly or by calling the `$PUTMSG` system service. It must either return a condition value in `R0` as a function value, or call `LIB$SIGNAL` or `LIB$STOP` to output all messages. (`LIB$SIGNAL` and `LIB$STOP` may be called either directly or indirectly.)
- A procedure should use device-independent services and procedures for input/output.

2.5.1 Terminal Input/Output

The methods available for performing input/output to the terminal include the following:

- Queue I/O Request system service (\$QIO)

Using a \$QIO to perform terminal I/O is very efficient; however, it is also the most difficult method to use. \$QIOs use device-dependent services and are the most difficult to use from high-level languages of all methods discussed here, because there are more steps involved and because the calling interface requires more knowledge from the caller than RMS services. Several steps may be required for your procedure to use a \$QIO. These may include constructing item lists, writing AST routines, assigning an I/O channel, queueing an I/O request, testing to ensure that the request was successfully queued and completed, and deassigning the I/O channel. (For further information on \$QIOs, see Section 7 of the *VMS System Services Reference Manual*.)

- VMS Record Management Services (RMS)

VMS RMS services are device-independent and provide general-purpose services that are easier to call than \$QIOs. However, it is often not convenient to construct the access control blocks (FAB, RAB, and so forth) required by VMS RMS services from a high-level language. (For further information on RMS services, see the *VMS Record Management Services Manual*.)

- Language I/O statements

Language I/O statements are provided for all high-level languages. These statements are easy to use and provide simple I/O and data formatting for the high-level language user. Native language I/O statements make it unnecessary for the high-level language user to call \$QIO or VMS RMS services directly; these calls are made by the compiled code on your behalf. However, low-level and medium-level languages (VAX MACRO and BLISS32) have no built-in language I/O statements and must use \$QIO and VMS RMS services for terminal and file I/O. (For further information, see the appropriate language reference manual.)

- Screen Management Procedures in the Run-Time Library (SMG\$)

SMG\$ procedures provide an easy-to-call interface for high-level languages. They are device-independent and aid in the composition of complex screen images. The SMG\$ facility in the Run-Time Library provides "screen composition operations; that is, SMG\$ makes it easy for an application to divide its screen into multiple regions and provides functions for manipulating those regions. Other features provided by SMG\$ procedures are as follows:

- Output to virtual displays
- Input from a virtual keyboard or locator device
- The ability to perform asynchronous input
- Built-in minimal screen updating
- Optional buffering and batching to optimize performance
- The ability to trap broadcast messages

Design

2.5 Using Input/Output

- The option of performing output to a file or a hardcopy device
- Support for foreign (non-DIGITAL) terminals
- Subprocess manipulation

For further information about SMG\$ procedures, see Section 3 of the *VMS Run-Time Library Routines Volume*.

During I/O to the terminal it is important that the procedure and the main program cooperate in controlling the terminal screen. For example, an I/O procedure may write something to the terminal screen that the calling program wants to erase. The calling program must know both what and where that information is, in order to erase it. The calling program and the called procedure must communicate by passing arguments that define which part of the screen will be accessed by each. The Run-Time Library contains Screen Management (SMG\$) procedures for this purpose.

It is essential that you do not combine different methods of I/O within your application. Problems can arise if the calling program and the called procedure use different methods of I/O. Each method of performing input/output maintains some knowledge of what is on the terminal screen. At the very least, the current cursor position is remembered. If another type of I/O is performed, that information is not updated and, therefore, becomes incorrect. The results of any subsequent I/O would be unpredictable. If you must combine other methods with uses of SMG\$ procedures, there are SMG\$ procedures that can be utilized to aid such an integration.

2.5.2 File Input/Output

File I/O can be performed by the following methods:

- Block I/O

Uses system services to map a section of the file to the process virtual address space. No notion of records.

- VMS Record Management Services (RMS)

VMS RMS provides a variety of file organizations and access modes from which you can select the processing techniques best suited to your application. VMS RMS supports the sequential, relative, and indexed-sequential file organizations. These modes allow you to access records within these files sequentially, randomly by key value or relative record number, or randomly by the record's file address (RFA). It is usually not necessary to call VMS RMS services directly from high-level languages. Consult your language reference manual for specific information on performing record management operations in the language you are using. (For further information on VMS RMS services, see the *VMS Record Management Services Manual*.)

- Language I/O

The compiled code in most high-level languages calls a Run-Time Library language support procedure for file operations. The Run-Time Library procedures normally call VMS RMS services. Thus, most VMS RMS features are available to the high-level languages user without calling RMS directly. Language I/O statements are suitable for either data files or output files. Low- and medium-level languages (VAX MACRO and BLISS32) do not have any language I/O statements and must call RMS

services directly. (For further information, see the appropriate language reference manual.)

2.6 Beginning the Internal Documentation

You must document each procedure carefully so you and others know what the procedure does. In most cases, a module should contain only one procedure.

2.6.1 How to Write a Module Description

You should place a preface containing the following information at the front of each module:

Title:	Gives the module name followed by a one-line functional description.
Version:	Gives the version and a three-digit edit number. Generally 1-001 is the original version.
Facility:	Gives a description of the library facility, such as general utility library (LIB).
Abstract:	Gives a short (three to six lines) functional description of the module.
Environment:	Lists any special environmental assumptions that the module can make. These include assumptions made at both compilation and execution time that could affect either the hardware or software environments. Describes situations that the module assumes during execution time and the optional elements of the VMS Modular Programming Standard that your module does not follow. Indicates the reentrancy characteristics of the procedures in this module. Each procedure is either fully-reentrant, AST-reentrant, or non-reentrant.
Author:	Includes your name and the date the module was created.
Modified by:	Includes the modification number, name of modifying programmer, modification date, and a list of the modifications.

This concludes the preface. End the preface with a page delimiter. The actual code for the module follows the preface.

Example 2-4 shows a sample module description.

Design

2.6 Beginning the Internal Documentation

Example 2-4 A Sample Module Description

```
PROGRAM GRA_CUBE                ! Create representation of a cube
!+
! VERSION:      1-002
!
! FACILITY:     User Graphics Computation Library
!
! ABSTRACT:     This module contains a procedure to create a mathematical
!               representation of a cube, GRA_CUBE.
!
! ENVIRONMENT:  User Mode, AST-reentrant
!
! AUTHOR:       Heather MacDonald      CREATION DATE:  14-Sep-1988
!
! MODIFIED BY:
! 1-001 - Original.  DWS 14-Sep-1988
! 1-002 - Fix a minor bug in cube volume computation.  MDL 15-Mar-1989
!-
```

2.6.2 How to Write a Procedure Description

You should place a procedure description at the beginning of each procedure in a module.

Always list each of the following topics regardless of whether or not they are actually present. For example, if a procedure has no implicit inputs, write the following:

```
!
!   Implicit Inputs:
!
!       NONE
!
```


2.6 Beginning the Internal Documentation

The procedure description includes the following elements:

Functional description:	<p>Describes a procedure's purpose and completely documents its interfaces.</p> <p>Includes the basis for any critical algorithms used, including literature references where applicable, and explains why a particular algorithm was chosen.</p> <p>Indicates the reentrancy characteristics of this procedure if they differ from those given in the module description.</p>
Calling sequence:	<p>Includes these elements in the following order:</p> <ol style="list-style-type: none"> 1 A return status, value argument, or CALL instruction 2 The procedure name 3 The argument list (typically a list of registers or arguments) <p>In MACRO, each argument is symbolically defined as the offset relative to the argument pointer (AP).</p> <p>Lists the arguments in the order they will appear in a high-level language. Each argument characteristic should also be included, using the procedure argument notation described in Appendix B.</p>
Formal arguments:	<p>Lists any explicit input, input-output, or output arguments. Includes a qualifying description with each argument. The arguments should be listed in the order they are listed in the calling sequence.</p>
Implicit inputs:	<p>Lists any inputs from storage, internal or external to the module, that are not specified in the argument list. Usually all that will appear here is "NONE". See Section 2.2.2.</p>
Implicit outputs:	<p>Lists any outputs to internal or external storage that are not specified in the argument list.</p>
Completion status or routine value:	<p>Lists the success or failure condition value symbols that could be returned as completion codes in R0. If your procedure returns a function value other than a condition value in R0, change the heading to "Routine value".</p>
Side Effects:	<p>Describes any functional side effects not evident from a procedure's calling sequence. This includes changes in storage allocation, process status, file operations, and possible signaled conditions. In general, you should document anything out of the ordinary that the procedure does to the environment. If a side effect modifies local or global storage locations, document it in the implicit output description instead.</p>

Example 2-5 shows a sample procedure description.

Design

2.6 Beginning the Internal Documentation

Example 2-5 A Sample Procedure Description

```
!++
! FUNCTIONAL DESCRIPTION:
!
!     Return the system date and time, using the caller's
!     semantics for his/her string.
!
!     Non-reentrant; uses static storage.
!
! FORMAL ARGUMENT(S):
!
!     RESULT_ADDR
!     VMS USAGE : char_string
!     TYPE      : character string
!     ACCESS    : write only
!     MECHANISM : by descriptor
!
!     Address of the descriptor into which the
!     system date and time is written.
!
! IMPLICIT INPUTS:
!
!     NONE
!
! IMPLICIT OUTPUTS:
!
!     NONE
!
! COMPLETION CODES:
!
!     SS$_NORMAL      Procedure successfully completed
!     LIB$_STRTRU     Success, but source string truncated
!
! SIDE EFFECTS:
!
!     Requests the current date and time from VMS.
!
!--
```

2.7 Planning for Signaling and Condition Handling

There are two methods available to a procedure for indicating to its caller whether it completed successfully. One method is to return a condition value in R0. The other method is to signal an error condition.

To provide a better user interface, all procedures in a facility should either return condition values or signal error conditions. Regardless of which method you choose, you should be consistent within the facility to make the procedures easier for the user to call.

2.7 Planning for Signaling and Condition Handling

2.7.1 Guidelines for Signaling Error Conditions

The signaling of an error condition is, in some instances, mandatory. Procedures that return a function value in R0 cannot return a condition value and therefore must signal any error conditions encountered.

However, to maintain efficiency, you might want other procedures to signal error conditions also. Checking the return status of a called procedure for repetitive calls can be time consuming and adversely affect the performance of the calling program. For example, if you are going to call a procedure 100 times within a loop and the chances of that procedure's failure are relatively small, you may not want to take the time to check the return status after each call to make sure that the condition value returned was `SS$_NORMAL`. Signaling error conditions is far more efficient in this type of application.

From the point of view of the calling program, handling a signaled condition is slightly more difficult than checking a returned condition value because it involves writing a condition handler to be invoked in the event that an error condition is signaled. However, handling a signaled condition allows the calling program to execute more efficiently.

To signal an error condition, your procedure uses either a condition handling mechanism provided by the source language, or it calls the Run-Time Library procedure `LIB$SIGNAL`. To use `LIB$SIGNAL`, your procedure calls `LIB$SIGNAL` and specifies the condition code and zero or more arguments specifying the environment of the condition. For further information about using `LIB$SIGNAL`, see Part II of the *VMS Run-Time Library Routines Volume*. The Run-Time Library also contains other procedures to help you with signaling and condition handling. These are described in Section 7, Table 7-1 of the *VMS Run-Time Library Routines Volume*.

2.7.2 Guidelines for Returning Condition Values

From the point of view of the calling program, it is considerably easier to check returned condition values than to handle signaled error conditions. When the condition value is being returned, the calling program does not need to include a condition handler. The calling program needs only to check the status of the returned value.

However, if you return condition values rather than signal error conditions, you return less information about the error condition to the calling program. It is recommended that you return condition values when the explanation of the error condition is simple and self-contained. For example, `LIB$GET_VM` returns a condition value, since the possible status conditions are self-contained and simple (for example, insufficient virtual memory).

According to the VAX Procedure Calling and Condition Handling Standard, the status returned must be a VAX condition value. (For further information, see the *Introduction to VMS System Routines*.)

Design

2.7 Planning for Signaling and Condition Handling

2.7.3 When to Signal or Return Condition Values

To some degree, whether you decide to signal an error condition or return a condition value depends on the language you are using for your procedure. In some high-level languages, it is very difficult to write a condition handler to be invoked in the event that an error condition is signaled. (For further information on condition handling in your language, consult the appropriate language reference manual.)

Regardless of which language you are using, there are general guidelines for when to return a condition value and when to signal an error condition.

You should signal an error condition in the following situations:

- Your procedure returns a value in R0 and cannot return a condition value.
- Your procedure must execute quickly and checking the return status of a condition value would be inefficient.
- Your procedure will be executed repetitively and, therefore, checking the condition value returned would adversely affect your procedure's performance.
- The amount of information you wish to return about the error condition cannot be contained in a condition value.
- A useful error message requires information that cannot be determined until run time. For example, the FDL\$PARSE procedure must tell you which line of the FDL file was the cause of an error. Since the line number of the line containing the error cannot be determined until run time, the signal mechanism is preferred.
- You want to execute a specific condition handler in the event that an error condition is signaled.

You should return a condition value in the following situations:

- You wish to keep the error-handling mechanism simple.
- The speed of the error checking mechanism is not of great concern.
- The total possible errors that may be returned is a small number and sufficient information about those errors can be contained in the condition value returned.
- The functions provided by the procedure are so general that the procedure will be used in various levels and environments.

3 Coding

This chapter describes how to code modular procedures. It has the following sections:

- Coding Guidelines
- Initializing Modular Procedures
- Using VMS System Services
- Invoking Optional User-Action Routines
- Writing AST-Reentrant Code

3.1 Coding Guidelines

The coding guidelines discussed in this section are drawn from the VMS Modular Programming Standard (Appendix A). These coding guidelines are of two types: *required* and *recommended*. You must follow the sections marked *required* to ensure that your application is modular. DIGITAL highly recommends that you adhere to the guidelines presented in the sections marked *recommended*. Following these additional rules will help you produce consistent, uniform applications.

3.1.1 Writing Position-Independent Code (Required)

A module is position independent when it can execute correctly anywhere in virtual memory. When programming in a VAX high level language, it is difficult not to write position-independent code because of the way the languages are designed. However, when programming in VAX MACRO and BLISS, certain combinations of addressing modes and the address of the operand result in code that is not position independent.

For more information on writing position-independent code in VAX MACRO, see Section 3.2.2 of the *VMS Linker Utility Manual*.

3.1.2 Adhering to the Naming Conventions

The following guidelines apply to the naming of facilities, procedures, files, modules, and program sections. You are required to follow these conventions when choosing names for modules, PSECTs, and status codes. DIGITAL highly recommends that you also adhere to the other naming conventions.

Coding

3.1 Coding Guidelines

- Add optional spaces.
- Insert block comments.

3.1.4.1 Using Symbols in Place of Numbers (Recommended)

Use symbols, not numbers, as much as possible. Because symbols are mnemonic, they will make your programs clearer and provide more information for cross-reference listings. It is always recommended that you define a symbol for a constant that is used a number of times. If the value for that symbol changes, you only have to change that value where it is defined, not in every place that it is used in the program. Thus, using symbols instead of numbers makes procedures easier to maintain. Furthermore, if you use a symbol instead of a number, it is easy to see each place that the symbol is referenced in the program. The cross-reference listing for a symbol lists each place that the symbol is referenced in the program.

Example 3-1 shows the information that can be obtained about a symbol in a cross-reference listing. The figure shows the contents of the LIGHT.LIS file created by the following command:

```
$ PASCAL/LIS/CROSS LIGHT.PAS
```

Note that the cross-reference listing shows where the symbols are declared and where they are referenced.

3.1.4.2 Using Uppercase and Lowercase Characters (Recommended)

You should use uppercase characters for all source code except comments. You should use both uppercase and lowercase characters for all comments. Comments that are complete sentences should start with a capital letter and end with a period.

All source code must be coded in uppercase characters if you want your applications to be independent of the operating system. Although VMS and its language implementations are insensitive to case, other operating systems and language implementations are not. (Note that VAX C is an exception; VAX C is case sensitive.) Applications coded in mixed or lowercase characters may not be easily transportable to other systems.

The PASCAL program shown in Example 3-2 illustrates the proper use of uppercase and lowercase characters.

3.1.4.3 Adding Optional Spaces (Recommended)

A single space should follow a comma (,) and precede and follow an equal sign (=). A single space should precede a left parenthesis (()) or a left square bracket ([]) (except in MACRO), but not a left angle bracket (<). A space should also follow an exclamation mark (!) or semicolon (;) to separate a comment from the source code. The arithmetic operators plus and minus (+ and -) should be surrounded by spaces in expressions.

The BASIC program in Example 3-3 shows the proper use of optional spaces.

Example 3-1 A Sample Cross-Reference Listing Showing the References to the Symbol SPEED_OF_LIGHT

```

MAX_SPEED
01                               Source Listing

0001          PROGRAM MAX_SPEED;
0002          {+}
0003          { This program shows the use of a symbol
0004          { instead of a number to improve clarity.
0005          { The program declares a symbol for the
0006          { speed of light, and then uses the symbol
0007          { to calculate the speed of the Starship
0008          { Lollipop and the Villain Vessel.
0009          {-}
0010
0011          CONST
0012          SPEED_OF_LIGHT = 3 * (10**8);
0013
0014          VAR
0015          LOLLIPOP_SPEED  : INTEGER;
0016          VILLAIN_SPEED  : INTEGER;
0017
0018          BEGIN
0019          LOLLIPOP_SPEED := SPEED_OF_LIGHT**8;
0020          VILLAIN_SPEED := SPEED_OF_LIGHT**4;
0021          END.
0022

MAX_SPEED
01                               Cross Reference Listing

LOLLIPOP_SPEED  VAR  INTEGER { IN PROGRAM MAX_SPEED }
                  15      19 =
INTEGER          TYPE
                  *      15      16
VILLAIN_SPEED    VAR  INTEGER { IN PROGRAM MAX_SPEED }
                  16      20 =
MAX_SPEED        PROG  [PSECT($CODE)]
                  1
SPEED_OF_LIGHT    CONST 300000000 { IN PROGRAM MAX_SPEED }
                  12      19      20

```

3.1.4.4 Inserting Block Comments (Recommended)

You can comment on blocks of statements by writing one or more lines preceding the block. The first comment line contains a single plus sign (+); the last comment line contains a single minus sign (-). Block comments do not need to be set off by additional blank lines; the two flag lines starting with a plus sign and minus sign serve that purpose. However, you may still wish to set off block comments with blank lines to improve readability. Comment delimiters are followed by one space, except when followed by the first plus sign and the last minus sign, as shown in Example 3-4.

Coding

3.1 Coding Guidelines

Example 3-2 PASCAL Program Showing Use of Uppercase and Lowercase Characters in Code

```
{+}
{ Program to call LIB$LP_LINES and determine the
{ number of lines per line printer page.
{-}
PROGRAM LINES(OUTPUT);
{+}
{ Declare the external procedure used by this
{ program.
{-}
FUNCTION LIB$LP_LINES : INTEGER; EXTERN;
{+}
{ Call LIB$LP_LINES and print the result.
{-}
BEGIN
    WRITELN('Each page contains ',LIB$LP_LINES,' lines.');
```

Example 3-3 BASIC Program Showing Use of Optional Spaces in Code

```
10      !+
        ! The following BASIC program converts a character
        ! string representing a hexadecimal value to a
        ! longword, then adds one to the result.
        !
        ! Declare the external routine used.
        !-
        EXTERNAL LONG FUNCTION OTS$CVT_TZ_L
        !+
        ! Perform the conversion.
        !-
        HEXVAL$ = "80012BFA"
        RET_STAT% = OTS$CVT_TZ_L (HEXVAL$, HEX% )
        !+
        ! Add one to the result.
        !-
        HEX% = HEX% + 1
        END
```

Example 3-4 FORTRAN Program Showing Use of Block Comments in Code

```
!+
!      This program demonstrates a call to the
!      Run-Time Library procedure STR$PREFIX.
!
!      Initialize the strings to be used.
!-
        A$ = "ABC"
        B$ = "DEF"
!+
!      Call STR$PREFIX
!-
        ISTAT = STR$PREFIX (A$, B$)
        END
```

3.1.5 Using VMS System Services

Not all VMS system services are modular, according to the definitions in this manual. Procedures that call nonmodular system services are nonmodular themselves. If your procedure uses a nonmodular system service, you should list the system service in the "Side Effects" section of the procedure description. (For information about the procedure description, see Section 2.5.2.) For further information about particular system services and modularity, see the *VMS System Services Reference Manual*.

3.1.6 Invoking Optional User Action Routines

An optional user action routine is a useful way to let the calling program gain control at a critical point in your procedure's algorithm. Success routines and error routines are the most common user action routines. Control is passed from your procedure to the optional error routine if the specified error is encountered within your procedure. To transfer control, the calling program must pass the user action routine as an argument to the called procedure. To make it easy for the calling program to pass information to its action routine, your procedure should supply an optional **user-arg** argument that the calling program can pass to its action routine. Your procedure merely copies the argument list entry of the user argument, if present, to the argument list it passes to the action routine. This achieves the same effect as up-level addressing.

There are two VAX data types that can be used to pass a user action routine as an argument: (1) the *procedure entry mask* and (2) the *bound procedure value*. To provide an interface to a user action routine for your procedure, you must first decide whether to use the procedure entry mask or the bound procedure value data type. The contents of the argument list entry for your action routine will differ, depending on whether you are specifying a procedure entry mask or a bound procedure value.

3.1.6.1 Procedure Entry Mask

The procedure entry mask (DSC\$K_DTYPE_ZEM) is the simplest to use. It is used by FORTRAN and most other languages and is expected by the Run-Time Library procedures, Record Management Services, and system services.

For a procedure entry mask passed by reference, the argument list entry contains the address of the procedure entry mask to be called. To provide a user action routine using the procedure entry mask, your procedure should have the following calling sequence:

```
CALL myproc [action-routine [,user-arg]]
```

In this example, **action-routine** is a function call of the procedure entry mask type that is passed by reference, and **user-arg** is unspecified.

Coding

3.1 Coding Guidelines

3.1.6.2 Bound Procedure Value

The bound procedure value (DSC\$K_DTYPE_BPV) is used by PASCAL and other languages where context of the procedure must be known. The procedure might do up-level addressing of a variable defined in a syntactically outer block and, hence, allocated in another frame. (If you use a procedure entry mask, this context is specified in the **user-arg** argument.)

For a bound procedure value passed by reference, the argument list entry contains the address of two longwords. The first longword contains the address of the procedure and the second contains the environment pointer to be loaded into R1 before the procedure is called. This environment pointer allows you to specify the context of your action routine enabling you to do up-level addressing. To provide a user action routine using the bound procedure value passed by reference, the calling sequence is as follows:

```
CALL myproc [action-routine [,user-arg]]
```

In this example, **action-routine** is a function call of the bound procedure value type that is passed by reference, and **user-arg** is unspecified.

If you want to use the bound procedure value data type to pass access to a user routine specified as a procedure entry mask, then you must pass the first longword by value and omit the second longword. Then, the user action routine would have this calling sequence:

```
status = action-routine (...[,user-arg])
```

In this example, status is a longword condition value that is passed by value, and **user-arg** is unspecified. Your procedure copies the 32-bit argument list entry passed by the calling program to the argument list provided to the action routine. Thus, the calling program and its action routine can communicate using any data type, access type, passing mechanism, or VMS usage.

3.2 Initializing Modular Procedures

Some modular procedures must initialize themselves before they can execute correctly. Following are examples of initialization:

- Storing in static storage a value that can only be determined at run time
- Declaring an exit handler using the \$DCLEXH system service
- Allocating a process-wide resource once
- Opening a file the first time the procedure is called

You must perform initialization carefully to maintain modularity.

Initialization must not affect the calling program. Therefore, avoid initializing by providing an entry point that must be called before any other entry point is called. Providing an entry point that must be called first forces the calling program to provide an initialization entry point to its caller, and so forth. Also, you would have to rewrite your calling programs if you needed to substitute a procedure with an initialization call for one without an initialization call.

3.2 Initializing Modular Procedures

If your procedure uses LIB\$INITIALIZE, you must preserve a modular environment that does not conflict with the environment set by any other procedure using LIB\$INITIALIZE. (See the *VMS Run-Time Library Routines Volume* for more information.)

Following are several ways to initialize a procedure:

- Initialize at compile or link time.
- Use the mechanism provided by LIB\$INITIALIZE to perform initialization once for each image activation.
- Set a first-time flag at run time.
- Initialize storage each time it is allocated at run time.
- Initialize storage each time a procedure is called at run time.

The use of each method is explained in the following sections. Figure 3-2 summarizes these methods.

Figure 3-2 Methods of Initializing

Initialization Needed	Method				
	Initialize at Compile/Link Time	Call LIB\$INITIALIZE Before Main Program (At Run Time)	Set a First Time Flag (At Run Time)	Initialize Each Time It Is Allocated (At Run Time)	Initialize Each Time Procedure Is Called (At Run Time)
Of Static Storage:	•	•	•		
Of Stack Storage:					•
Of Heap Storage:				•	
To Allocate Resources:		•	•		
To Set Up \$EXIT Handler:		•	•		
To Open a Process-Permanent File:		•	•		
To Set Up a Handler Before the Main Program:		•			

ZK-3085-84

Coding

3.2 Initializing Modular Procedures

3.2.1 Initializing Storage

For a procedure to produce predictable results, all statically and dynamically allocated areas must be initialized to known values before they are read. Initialization of dynamically allocated stack and heap data involves writing the data after each allocation and before reading it.

If your procedure has static storage, it is usually initialized to zero. In some languages (for example BLISS), you do not need to explicitly initialize static store. These languages will automatically initialize static storage to zero. To see if the language you are using initializes static storage implicitly, refer to your reference manual for that language.

There are three ways to explicitly initialize storage: you can use an initialization statement, test and set a first-time flag at run time, or use LIB\$INITIALIZE. The method of testing and setting a first-time flag is explained in Section 3.3.4.2.

Figure 3-3 provides examples in the major VAX languages of how to initialize a longword, DAT, in static storage using an initialization statement.

3.2.2 Testing and Setting a First-Time Flag

To do first-time initialization, your procedure can test and then set to one a statically allocated first-time flag each time it is called. This flag is initialized to zero at compile or link time. Setting and testing the flag with the VAX instruction Branch on Bit Set and Set (BBSS) insures that initialization is executed exactly once. (You can also use BBSSI or the Run-Time Library procedure LIB\$BBSSI.)

However, if your implementation language does not have access to VAX instructions and the procedure is to be AST-reentrant, it must follow these steps:

- 1 Test the first-time flag.
- 2 If the first-time flag is set, initialization is complete.
- 3 If the first-time flag is not set, disable ASTs. Remember the previous state of AST enable, and retest the flag.
- 4 If the first-time flag is now set, then initialization was performed by an AST that occurred between the first test and the AST disable; enable ASTs if remembered state of ASTs was enable. Initialization is now complete.
- 5 If the first-time flag is not set, perform the initialization.
- 6 Set the flag.
- 7 Enable ASTs if remembered state of ASTs was enable — initialization is complete.

For additional information, see Section 3.3.

Note: ASTs should be enabled in step 4 or step 7 only if they were enabled before step 3. The \$SETAST system service, used to disable ASTs, indicates whether ASTs were enabled when the procedure was called.

Coding

3.2 Initializing Modular Procedures

Figure 3-3 How to Initialize Static Storage

Language	Statement	Initialized Value	Time of Initialization
Ada	x : INTEGER := 1	1	Elaboration time
BASIC	BASIC does not permit static storage within a module, only common static storage.		
BLISS	OWN DAT;	0	Compile time
	OWN DAT INITIAL(0);	0	Compile time
	OWN DAT INITIAL(100);	100	Compile time
C	STATIC INTEGER x;	0	Compile time
	STATIC INTEGER x = 1	1	Compile time
	EXTERNAL INTEGER x; Defined externally ¹		Compile time
	INTEGER x;	0	Compile time
	INTEGER x = 1	1	Compile time
	GLOBALDEF INTEGER x;	0	Compile time
	GLOBALDEF INTEGER x = 1	1	Compile time
COBOL	01 NUM PIC	0	Compile time
	01 NUM PIC 9 VALUE 0.	0	Compile time
	01 NUM PIC 9(3) VALUE 100.	100	Compile time
DIBOL	At compile time, fields within records, commons, and/or groups are initialized to spaces or zeros (depending on data type).		
FORTRAN	INTEGER*4 DAT	0	Compile time
	INTEGER*4 DAT/0/	0	Compile time
	DATA DAT /0/	0	Compile time
	DATA DAT /100/	100	Compile time
MACRO	DAT: .BKLL 1	0	Compile time
	DAT: .LONG 0	0	Compile time
	DAT: .LONG 100	100	Compile time
PASCAL VAR			
	DAT : [STATIC] INTEGER;	0	Compile time
	DAT : [STATIC] INTEGER := 0;	0	Compile time
	DAT : INTEGER := 100;	100	Compile time
PL/I	STATIC INIT(2)	2	Compile time
	EXTERNAL INIT(3)	3	Compile time
	GLOBALDEF INIT(4)	4	Compile time
	GLOBALREF INIT(5)	5	Compile time
RPG	RPG II has static storage at the module level only. Numeric variables are initialized to zero and alphanumeric variables are initialized to spaces at compile time.		
SCAN	No initialization clauses—use assignment statement.		

¹ You cannot initialize a variable declared with an external attribute.

Coding

3.2 Initializing Modular Procedures

Example 3-5 illustrates the use of a first-time flag in a PASCAL program to allocate a resource.

Example 3-5 PASCAL Program Which Uses a First-Time Flag

```
{+}
{ Program to demonstrate the use of a first-time flag when allocating
{ a resource. This technique is AST-reentrant, but is NOT multi-thread
{ reentrant.
{-}

PROGRAM ALLOCATE;

CONST
    VM_SIZE = 512;

VAR
    INITIALIZED : BOOLEAN := FALSE;
    VM_ADDRESS : INTEGER := 0;
    AST_STATUS : INTEGER := 0;
    VM_STATUS : INTEGER := 0;
    DISABLE : INTEGER := 0;

FUNCTION LIB$GET_VM (SIZE : INTEGER; VAR ADDR : INTEGER) : INTEGER; EXTERNAL;
FUNCTION SYS$SETAST (VAR STATUS : INTEGER) : INTEGER; EXTERNAL;

BEGIN
    {+}
    { Check the first-time flag. If set, initialization has been
    { performed already.
    {-}

    IF NOT (INITIALIZED)
    THEN
        BEGIN
            {+}
            { Disable AST's, and remember the previous state.
            {-}

            AST_STATUS := SYS$SETAST (DISABLE);

            {+}
            { Now, recheck the flag. If it is now set, initialization was
            { performed by another invocation of this procedure between when
            { the flag was first tested and now. Otherwise, initialization
            { is performed here.
            {-}

            IF NOT (INITIALIZED)
            THEN
                BEGIN
                    {+}
                    { Perform the initialization.
                    {-}

                    VM_STATUS := LIB$GET_VM (VM_SIZE, VM_ADDRESS);

                    {+}
                    { Set the first-time flag, indicating initialization complete.
                    {-}
```

Example 3-5 Cont'd. on next page

Example 3-5 (Cont.) PASCAL Program Which Uses a First-Time Flag

```

        INITIALIZED := TRUE;
        END;

    {+}
    { Restore AST's to the previous state.
    {-}

    AST_STATUS := SYS$SETAST (AST_STATUS);
    END;
END.
```

3.2.3 Using LIB\$INITIALIZE

One of the ways that you can initialize a value at run time is by using the Run-Time Library procedure LIB\$INITIALIZE. An example of a value that you may need to initialize at run time is a seed for a random number generator.

The following six steps are needed to use LIB\$INITIALIZE to initialize a value at run time:

- 1 Write the main program.
- 2 Write an initialization procedure.
- 3 Write a MACRO or BLISS program to add the address of that initialization procedure to PSECT LIB\$INITIALIZE.
- 4 Compile the initialization procedure, main program, and MACRO program.
- 5 Link the initialization procedure, main program, and MACRO program.
- 6 Run the main program.

Assuming that you have completed the main program, the first thing that you need to do is to write an initialization procedure. If, for example, you were going to use LIB\$INITIALIZE to initialize a value for a random number generator, you might write an initialization procedure to set the seed equal to the current time. This would generate a different seed for each initialization because the time is constantly changing. One possible initialization procedure is shown in Example 3-6.

Once you have defined the initialization procedure, you must write the MACRO program to add the address of that initialization procedure to PSECT LIB\$INITIALIZE. The format for this MACRO program is very simple, as seen in Example 3-7.

To modify this MACRO program for use in your own procedures, substitute the name of your initialization procedure for MY_INIT_ROUTINE.

Coding

3.2 Initializing Modular Procedures

Example 3-6 BASIC Initialization Procedure for LIB\$INITIALIZE

```
100      !+
        ! Initialization routine. A common piece of data, called SEED,
        ! is initialized based on the number of CPU seconds used by
        ! this process so far.
        !-
        SUB MY_INIT_ROUTINE(ONE,TWO,THREE,FOUR,FIVE,SIX)
        COMMON (MY_DATA) LONG SEED
        PRINT "Now in initialization routine."
        CURRENT_TIME = TIME(1)
        SEED = CURRENT_TIME
        END SUB
```

Example 3-7 Program to Add Address to PSECT LIB\$INITIALIZE

```
;+
; Make references to external routines used.
;-
        .EXTRN LIB$INITIALIZE
        .EXTRN MY_INIT_ROUTINE
;+
; Make a contribution to the PSECT LIB$INITIALIZE.
;-
        .PSECT LIB$INITIALIZE USR,GBL,NOEXE,NOWRT, LONG
        .ADDRESS MY_INIT_ROUTINE
        .END
```

Once you have written the initialization procedure and the MACRO program to add the dispatch address to PSECT LIB\$INITIALIZE, you can link and run your program. The sample program in Example 3-8 can be initialized in this manner.

Example 3-8 BASIC Main Program

```
10      !+
        ! Mainline. The value of SEED is printed.
        ! The linker initializes this value to zero, but because LIB$INITIALIZE
        ! is used, an initialization routine is run before control is transferred
        ! here, and the value of SEED is changed to a somewhat random value.
        !-
        COMMON (MY_DATA) LONG SEED
        PRINT "Now in mainline. The seed is initialized to: ";SEED
32767   END
```

To run LIB\$INITIALIZE on the program in Example 3-8 and thus initialize the value of SEED at run-time, enter the following commands:

```
$ BASIC MAIN
$ BASIC INIT
$ MACRO LIBRARY
$ LINK MAIN,INIT,LIBRARY
$ RUN MAIN
```

The following is an example of the output generated by these steps:

```
Now in initialization routine.
Now in mainline. The seed is initialized to: 4099
```


If your procedure establishes a condition handler by calling `LIB$INITIALIZE` before a main program, the action of this handler might conflict with other condition handlers established by other procedures before the main program.

3.3 Writing AST-Reentrant Code

This section describes coding techniques for modular procedures that use the VMS asynchronous system trap (AST) interrupt mechanism or that permit calling programs to use it.

All modular procedures should be AST-reentrant so they can be called from any program. If your procedure is not AST-reentrant or calls any procedure that is not, your program documentation should specify this to warn others against using your procedure.

3.3.1 What is an AST?

An asynchronous system trap (AST) is a VMS mechanism for providing a software interrupt when an external event occurs. One example of this type of interrupt occurs when the user presses CTRL/C. When the external event occurs, the VMS operating system interrupts the execution of the current process and calls a procedure that you supply. This procedure is what we refer to as the AST handler.

Some VMS system services let an external event interrupt a process. Because the interrupt occurs out of sequence with respect to process execution, the interrupt mechanism is called an "asynchronous" system trap. The AST interrupt transfers control to the AST handler that services the event. This AST handler can call other procedures, including library procedures.

The AST handler you provide and any procedures it calls are said to be executing at AST level. While at AST level a process cannot be interrupted a second time at the same access mode. The process runs to completion at the AST level before the non-AST level procedure resumes.

A process is executing either at AST level or at non-AST level and thus consists of two "threads of execution," one thread at each level. Keep in mind that these levels are threads of the same process and not separate processes.

When your AST handler finishes servicing the event, it returns control to its caller. The interrupted procedure continues execution from the point of interruption.

For example, you could call the Set Timer system service (`$SETIMR`) to specify the address of an AST-level procedure to be executed after a specified amount of time has elapsed. At the specified time, the system generates an AST interrupt by stopping the procedure that is currently executing and calling the specified AST handler.

For information on the implementation of AST interrupts by system services, see the *VMS System Services Reference Manual*.

3.3.2 AST-Reentrancy Versus Full-Reentrancy

A procedure is AST-reentrant if it meets the following conditions:

- It can be interrupted at any point, permitting itself or any related procedure to be called (reentered).
- It executes correctly when it continues from the point of interruption.

Do not confuse the term *AST-reentrant* with the term *fully-reentrant*. Full reentrancy refers to a more restrictive set of conditions.

In an AST-reentrant environment, the AST thread is expected to complete regardless of whether it encounters a locked resource. When the AST thread encounters a locked resource in an AST-reentrant environment, it expects to be given a "new" resource, or else it is expected to return an error message. It is never expected to wait for the resource that the non-AST level has locked.

In a fully-reentrant environment, all threads are treated equally when they encounter a locked resource; they wait for the resource to be freed. In a fully-reentrant environment, AST threads are not given any special treatment. The VAX Ada environment is an example of a fully-reentrant environment. In such a situation, there can be more than two threads of concurrent execution, and each thread can alternately progress toward an end.

Note: It is highly desirable that future code satisfy the more stringent requirement of being fully-reentrant. Full reentrancy is important for procedures that will be called from multi-thread environments, such as Ada tasks. For further information, refer to the VAX Ada documentation.

3.3.3 Guidelines for Writing AST-Reentrant Modular Procedures

To use AST interrupts, you must write an AST handler to take control at AST level. An AST handler can be written in any language. Because the particulars of writing an AST handler differ from one language to the next, see the reference manual for the language you are working in for more details.

In general, an AST handler must follow these guidelines:

- It must be separate from the procedure that is currently executing.
- It must not modify data or instructions used by the interrupted procedure or its callers.
- The AST handler must be callable with a CALLG or CALLS instruction and must return with a RET instruction.
- If it modifies any registers other than R0 and R1, it must set bits in the entry mask to save the contents of the registers.
- If it calls any other procedures, they must all be AST-reentrant.
- The AST handler cannot stall or use "busy wait" to avoid being called before the non-AST level is out of a critical section of code. Once the AST handler has begun executing, it cannot be interrupted by anything at a non-AST level. In fact, the only thing that can interrupt the AST handler is another procedure running at AST level in a different access mode.

3.3 Writing AST-Reentrant Code

If you attempt to use a "busy wait" and expect to change the condition from the non-AST level, the AST level circles the "busy wait" in an infinite loop. The process continues to loop because the non-AST level does not continue executing until the AST thread has finished and thus is never able to change the value in the "busy wait" condition.

- You cannot use the lock manager to protect a resource being accessed at non-AST level from being accessed at AST level. The lock manager is designed to lock resources between separate processes, not different threads (AST and non-AST) of the same process.
- Avoid using static storage. A procedure that does not use static storage, calls only AST-reentrant procedures, and does no up-level addressing is automatically AST-reentrant.

3.3.4 How to Eliminate Race Conditions During Concurrent Access

There are two problems that you might encounter in using AST interrupts: race conditions and deadlocks. A race condition occurs when your AST handler attempts to use a nonshareable resource already in use by the non-AST thread of execution.

If you allow the AST handler to wait for the resource (for example, by waiting for an event flag to be set by the non-AST level code of the same access mode), you have caused a form of deadlock. The deadlock occurs because the non-AST level code cannot execute to free the resource until the AST-level code has finished executing. The AST level code cannot continue either, because the non-AST level code has effectively locked the resource.

A race condition occurs when you attempt to access or modify the same data in static storage by both the AST and non-AST level of a process. For example, if an AST begins executing while the non-AST level is modifying data in static storage, that data may be left in a non-stable state while the AST handler executes. To prevent a race condition, you should allow only one thread at a time to modify data.

If a procedure does not modify any static storage, then it is both AST-reentrant and fully-reentrant. Your procedure can eliminate conflict when accessing and modifying data in static storage in the following ways:

- Performing all accessing or modification in a single uninterruptable instruction. (This technique also makes the procedure fully-reentrant.)
- Detecting concurrency of access to data using "test and set" instructions at entry to and exit from data storage. The procedure may then report an error, or retry the operation (when appropriate) if concurrency is detected.
- Keeping a call-in-progress count that is incremented when your procedure is called and decremented when it returns. The count is used as an index into separate allocated areas.
- Disabling AST interrupts upon entry and restore the enable state upon exiting.

Coding

3.3 Writing AST-Reentrant Code

3.3.4.1

Performing All Accesses in One Instruction

All data modification in static storage can be performed in a single uninterruptable instruction for some applications. However, this method applies only to MACRO, and even there does not apply to emulated instructions.

For example, you can use queue instructions to maintain a linked list in a single instruction instead of modifying the forward and backward fields of the list in several instructions. You can use a single queue instruction at the beginning of your procedure to remove one section, and another can be used at the end to insert the section back in the queue.

While a section is removed from the queue, your procedure can modify data in it. If an AST interrupt occurs while the section is removed, a different section of data is used instead, thus avoiding conflicts with the interrupted procedure.

Example 3-9 illustrates an AST-reentrant procedure that uses queue instructions to control allocation of quadword blocks.

Example 3-9 MACRO Program Showing Use of Queue Instructions to Perform All Accesses in a Single Instruction

```
.PSECT _LIB_DATA PIC,USR,CON,REL,LCL,NOSHR,NOEXE,RD,WRT
FLAG: .LONG 0 ; First-time flag1
Q_HED .LONG 0,0
.PSECT _LIB_CODE PIC,USR,CON,REL,LCL,SHR,EXE,RD,NOWRT
.ENTRY LIB_GET_X,~M<>
BBC FLAG, FIRST ; Branch on 1st call only
TRY: REMQUE @Q_HED, RO ; RO = address of queue
BVS 10$ ; Branch if empty and fill
RET

10$: BSBB FILL ; Fill queues
BRB TRY ; Try again

;+
; Here on first call only
;-
FIRST: $SETAST #0 ; Disable ASTs, RO=old setting
BBSS FLAG, 20$ ; Branch if already set
MOVAL Q_HED, Q_HED ; Make queue empty
MOVAL Q_HED, Q_HED+4 ; Back pointer too
BSBB FILL ; Fill queues
20$: CMPL #SS$_WASSET, RO ; were ASTs enabled before?
BNEQ TRY ; No, leave disabled, retry
$SETAST 1 ; Yes, enable ASTs
BRB TRY ; Try again

FILL: get space for 10 quadwords by calling LIB$GET_VM
and insert in queue using INSQUE
RSB
```

¹ This example could be recoded using REMQHI and INSQHI and avoid the need for a first time flag.

3.3.4.2 Using "Test and Set" Instructions

One method of eliminating the possibility of a race condition or deadlock is to use "test and set" instructions to detect concurrent access. You can detect concurrent access of static storage at both AST and non-AST levels by adding the following steps to your procedures:

- Place a branch on bit set and then set BBSS or BBSSI instruction immediately before your procedure accesses static storage.
- Access or modify static storage, or both.
- Place a branch on bit clear and then clear BBCC or BBCCI instruction immediately after your procedure has completed access to static storage.

The BBSS instruction detects that a concurrency conflict is about to take place before static storage has been accessed. If the storage is being accessed by multiple processors, you must use BBSSI and BBCCI.

There are two alternate techniques for resolving concurrency conflicts detected by the BBSS and BBCC instructions:

- Use separate, statically allocated areas for storage at the AST and non-AST levels. When the BBSS instruction detects concurrency at the beginning, use the second allocated area. Note that this technique does not work if an exception condition occurs between execution of the BBSS instruction and the BBCC instruction or if your procedure has not established a condition handler. This is because a condition handler established by the calling program might also simultaneously call your procedure.
- Reexecute your procedure if concurrency is detected. When the BBCC instruction detects this concurrency, branch back to the beginning of your procedure and try again.

Example 3-10 illustrates the latter technique. This MACRO procedure, LIB_GET_INUM, allocates and deallocates identifying numbers.

Example 3-10 MACRO Program Showing Use of Test and Set Instructions

```
.TITLE LIB_GET_INUM  -- Allocate and deallocate id. nos. 1 - 10
TAB:  .WORD  0          ; Bitmap for flags
      .ENTRY  LIB_GET_INUM, ^M<>
10$:  FFC      #1, #10, TAB, RO    ; Find first free id, no.
      BEQ      20$                ; Branch if none free
      BBSS     RO, TAB, 10$        ; Indicate id. no. in use
      MOVL     RO, @4(AP)          ; Return id. no. found
      MOVL     #1, RO              ; Indicate success
      RET
20$:  CLRL     @4(AP)              ; Return 0
      CLRL     RO                  ; Indicate failure
      RET
      .END
```

Coding

3.3 Writing AST-Reentrant Code

3.3.4.3 Keeping a Call-in-Progress Count

If the database is to be kept separate between calls, you can keep track of when your procedure is called by using a call-in-progress count. Before database access, the count is incremented and used as an index for an address table of the separate databases. You should check for a count that exceeds the table length. After the database has been accessed, the count is decremented.

This technique has an advantage over the BBxx technique because it can handle more than two levels of reentrance. However, it is less reliable because an exception can cause the count never to be decremented, leading to an eventual procedure malfunction. You can avoid this by establishing a condition handler in your procedure.

3.3.4.4 Disabling AST Interrupts

A procedure is also considered AST-reentrant if AST interrupts are disabled while critical sections of code execute. However, this method of maintaining AST reentrancy is not recommended.

Sometimes the only way to avoid race conditions is to disable AST interrupts during the access to static storage and restore the state of the AST enable once the critical section of code has finished executing. However, this technique could adversely affect performance of real-time programs using AST interrupts. The \$SETAST system service, which is used to enable and disable AST interrupts, is time consuming. Therefore, you should avoid disabling AST interrupts whenever you can by using the techniques described in Section 3.3.4.1 to Section 3.3.4.3.

Try to minimize the number of instructions during which the AST interrupts are disabled. Before disabling AST interrupts, establish a condition handler to restore the AST level in case an exception or stack unwind occurs.

Example 3-11 demonstrates how \$SETAST can be used to disable ASTs and then restore the previous state of the enable.

Example 3-11 A FORTRAN Program Disabling and Restoring ASTs

```
!+
! This program demonstrates using the System
! Service SYS$SETAST to disable and then
! re-enable AST-interrupts.
!-
      INCLUDE '($SDEF)'
      INTEGER*4 SYS$SETAST

!+
! Turn off ASTs and remember the previous setting.
!-
      ISTAT = SYS$SETAST (%VAL(0))

!+
! The statements in the program during whose
! execution you want ASTs disabled.
!
! If ASTs were previously enabled,
! re-enable them.
!-
      IF (ISTAT .EQ. SS$_WASSET) CALL SYS$SETAST( %VAL(1))
      END
```


3.3.5 Performing Input/Output at AST-Level

If your procedure performs I/O using VMS RMS system services, there are several coding techniques you must follow for your procedure to be AST-reentrant:

- When opening process-permanent files — such as SYS\$INPUT, SYS\$OUTPUT, SYS\$COMMAND, or SYS\$ERROR — check for the VMS RMS error status RMS\$_ACT (active) after each \$CREATE or \$OPEN service. This error indicates that a record operation has already started for the process-permanent file. The error does not occur for files that are not process permanent, and the \$OPEN service follows the constraints of shared access to the file that may have been imposed by a previous \$OPEN service. If the error occurs, perform a \$WAIT using the same file access block (FAB). When control returns to your procedure, try the \$CREATE or \$OPEN service again. Repeat this sequence until it succeeds.
- When performing record I/O to any type of file, check for the RMS error status RMS\$_RSA (record stream active) or RMS\$_BUSY (structure in use) after each \$GET and \$PUT service. This error indicates that a record operation has already been started for the file. If the error occurs, perform a \$WAIT using the same record access block (RAB). When control returns to your procedure, try the \$GET or \$PUT service again. Repeat this procedure until it succeeds.
- Avoid storing data in a record access block (RAB) that VMS RMS could still be accessing. You can avoid this situation by doing either of the following:
 - Allocate the RAB on the stack so the AST and non-AST level have separate RABs.
 - Allocate RAB in heap or static storage along with a busy bit. The busy bit is tested and set using a BBSS instruction before the RAB is accessed. If the RAB is already busy, your procedure executes a \$WAIT using that RAB.

For synchronous input/output (I/O that is always completed before returning control to your procedure), you can allocate the RAB in either of these ways. However, the first method is more reliable, since it doesn't use static storage and therefore does not become corrupted if an exception is signaled.

For asynchronous I/O (when control is returned to your procedure before I/O is completed), you must use the second technique.

Coding

3.3 Writing AST-Reentrant Code

3.3.6 Condition Handling at AST Level

You should not allow an exception to propagate out of an AST handler since the exception might be caught by any procedure that is active at the time of the AST. Condition handlers for other active procedures might react as if the exception was caused by a procedure that they had called.

Another reason for not allowing exceptions to propagate out of an AST handler is that, for run-time environments that use multiple threads in a process such as VAX Ada, it cannot be determined which stack of the threads of execution is used to deliver the AST. (The AST is delivered on the stack of whichever thread is active at the time of the AST interrupt.)

It is best to catch all exceptions in the AST handler and not allow them to propagate.

4 Testing

A successful test system is one that uncovers errors. To facilitate successful testing you should begin planning for the testing phase during the design phase. Preliminary testing should start while you are coding the procedure. In this respect, some parts of the testing stage are concurrent with the coding stage in the software life cycle.

The two primary things that you should be testing for are as follows:

- 1 To make sure that the procedure you developed fulfills your requirements or specifications.

You must carefully test each aspect of functionality to ensure that your procedure does everything that it was intended to do and nothing that it was not intended to do. The methods you use to test this aspect of your procedure depend upon the functions your procedure performs.

- 2 To make sure that the procedure is modular and executes without error.

This chapter focuses on testing for modularity.

Modularity is especially important to procedures that are to be included in a library facility. If your procedure is in any way non-modular, it can adversely affect the results and performance of other procedures that call it. It is essential that your procedure completely adheres to the required sections of the VMS Modular Programming Standard contained in Appendix A.

To ensure modularity within your procedures, you should perform at least three types of tests:

- 1 Unit testing
- 2 Language-independence testing
- 3 Integration testing

Methods for designing and administering these types of tests are discussed in the following sections. Reentrancy and performance analysis are also discussed.

4.1 Unit Testing

Before you begin combining units of code (such as subprograms, subroutines, and internal procedures) to form your new procedure, it is essential to ensure that each of these units works separately. Thorough unit testing is important for the following reasons:

- Testing small units separately decreases the level of complexity within the tests.
- It is easier and faster to debug a small unit of code than it is to find an error within several units and their interfaces.
- It makes the integration stage that follows much easier if each of the separate units has been thoroughly tested and the problems corrected.

Testing

4.1 Unit Testing

- The earlier an error is found in development, the less expensive it is to fix.

There are three steps involved in unit testing:

- 1 Review the goals of your procedure.
- 2 Choose test cases.
- 3 Run the tests.

The goals of your procedure are chosen at the requirements or specifications stage. As we mentioned earlier, that stage is not discussed within this manual because it does not have a significant effect on modularity. However, it does have a significant effect upon whether your final product can be considered successful. If your product does not perform the functions or meet the requirements decided upon at the requirements or specifications stage, it is not a successful project. You should have at least one test for each of the requirements that your procedure was designed to fulfill.

You use the following two types of tests:

- Black box tests
- White box tests

Black box tests assume that you know nothing about the internal workings of the procedure that you are testing. All that you are interested in is the output that you receive for given sets of input.

White box tests (also called clear box tests) are more complicated because they are designed to step through particular sections of code or algorithms internal to the procedure. They assume that you know, in great detail, the internal workings of the procedure being tested.

4.1.1 Black Box Testing

When you are performing black box testing, you are interested only in the output you receive for particular input values. You simply execute the procedure repetitively using input from different classes. The best way to do this is to write a command procedure to execute the procedure a given number of times using test data that you supply. (For information on writing command procedures, see the *Guide to Using VMS Command Procedures*.)

You should execute your procedure with test cases from each of the following categories:

- Expected inputs

These are the values that you expect your procedure to receive most of the time.

- Boundary values

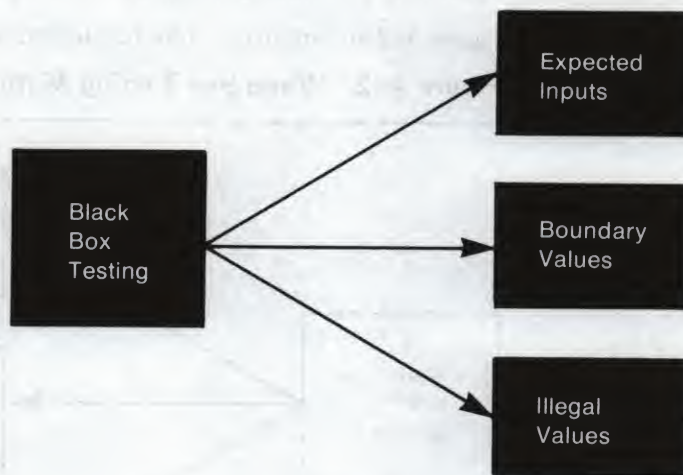
If your procedure expects an input value from 1 to 999, use 1 and 999 as test cases to make sure that your procedure returns the expected results for the boundary cases.

- Illegal values

Using the example above, what happens if your procedure receives as input a value that is less than 1 or greater than 999? Does the user receive a useful error message? Does the procedure simply stop, or does it attempt to use values outside of its limitations and simply return an incorrect answer? It is essential that you run the procedure using illegal input values to determine the answers to these questions.

Figure 4-1 summarizes the methods of black box testing.

Figure 4-1 Black Box Testing Methods



ZK-4071-85

4.1.2 White Box Testing

When performing white box testing, unlike black box testing, you must understand fully the internal workings of the procedure. Those internal workings, those specific lines of code, are in fact what you are testing.

The following steps are involved in white box testing:

1 Test each statement.

For this step, you need to provide sets of test values that ensure that every statement in the procedure is executed at least once. This includes all statements — even those executed only when optional arguments, user-supplied arguments, subroutines, user-action routines, or specific error codes are present.

2 Test each decision.

At this step, your goal is to provide test cases that ensure that each branch of a decision is executed at least once. In the case of a standard boolean decision, this generally requires providing two values; however, this number may be much greater in the case of compound or nested decisions.

3 Test each condition.

Testing

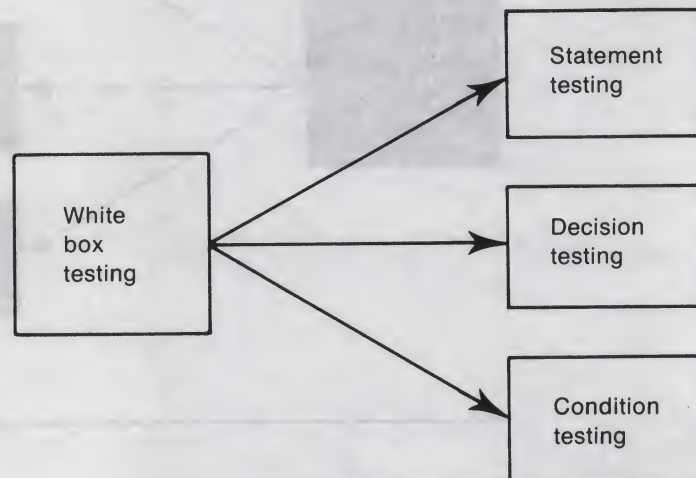
4.1 Unit Testing

Condition testing requires writing test cases that ensure each condition in a decision takes all possible outcomes at least once and each point of entry to the program or subroutine is invoked at least once. Multiple test values must be supplied in cases of compound and nested loops. In testing the entry points, remember to invoke any optional routines (either internal or external), as well as error handlers. If your procedure contains a JSB entry point, that entry point should also be tested.

Note that each type of white box testing finds a specific type of error. For example, statement testing does not find an error on a negative value for a condition if the statement is given a positive input the only time it is executed. Therefore, you must perform all three types of white box testing.

Figure 4-2 summarizes white box testing methods.

Figure 4-2 White Box Testing Methods



ZK-4069-85

4.2 Language-Independence Testing

For your procedures to be as useful as possible, they must be able to be called by programs in any language. Providing for language independence is essential to producing a useful procedure.

Testing for language independence is a very specific type of unit testing. It ensures that your program executes correctly regardless of the language from which it is called.

To test your procedures for language independence, write several driver programs in languages you have chosen randomly. The driver program need only contain a call to the procedure being tested.

If you do find that your procedures are not language independent, make sure that they conform to the following rules:

- All atomic data must be passed by reference and all strings must be passed by descriptor.

Testing

4.2 Language-Independence Testing

Adherence to this single guideline is the most important factor in achieving language independence.

- Statements that assume a particular language environment are NOT allowed.

For example, the statement `ON ERROR GO BACK` in a BASIC procedure assumes that the calling program is also written in BASIC.

4.3 Integration Testing

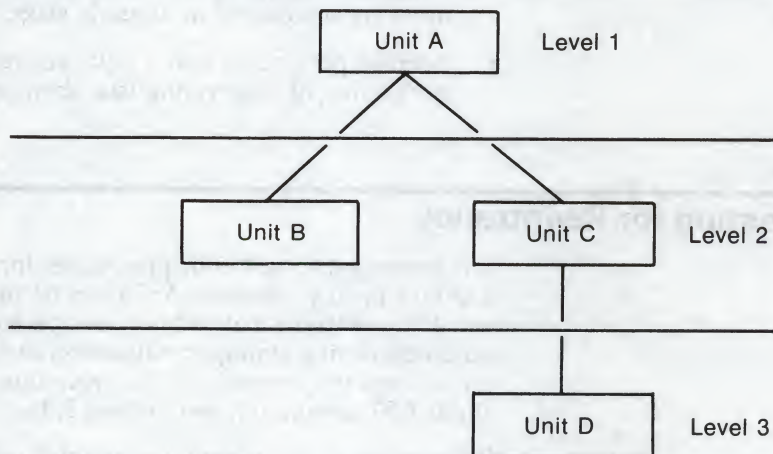
Integration testing is the next logical step following unit testing. Unit testing is designed to test each separate component. Depending on your procedure, that component might be a module, a subprogram, a subroutine, an internal procedure (face—name), or a particularly intrinsic piece of code. Once you have determined that each unit works separately, you need to determine that the units also work together to form the complete procedure.

Integration testing can be completed by either of two methods.

4.3.1 The "All at Once" Approach to Integration Testing

One method of integration testing is the "all at once" approach. In this method, you simply finish all of the units, link them together and test the completed structure all at once. Use of this method is strongly discouraged, because it makes it very difficult to find the location of errors. For example, look at the organization of the units in the sample procedure shown in Figure 4-3. Assume that we tested this procedure using the "all at once" approach and found an error; the procedure simply did not work. We would have no way of knowing whether the error was in unit A, unit B, unit C, or unit D.

Figure 4-3 A Sample Procedure for Integration Testing



ZK-4070-85

Testing

4.3 Integration Testing

4.3.2 The Incremental Approach to Integration Testing

The recommended approach to integration testing is called *incremental* testing. Incremental testing involves testing the procedure by starting with one unit and building on to it one unit at a time. Each unit should always be subjected to thorough unit testing before it is included in the integration tests.

Incremental integration testing is especially useful for finding the following types of error:

- Problems with the calling interface between units (for example, inconsistent ordering of arguments between the calling and called unit)
- Incorrect assumptions about what values are returned and the units to which they are returned
- Unexpected transfer of control between units

Using the sample procedure in Figure 4-3, complete the test of unit A on level 1 before proceeding to level 2 where you test units A and B in combination. At each level you correct any errors before proceeding to the next level. When you have completed the last step, you know that the entire procedure works correctly.

Because you started at the top of the sample procedure and added units incrementally from lower levels, you were using the *top-down* approach to integration testing. You could just as easily have started at Level 3 and used the *bottom-up* approach.

As you can see from the example, there are several distinct advantages to incremental integration testing:

- It is not necessary to wait until the procedure is complete to begin integration testing.
- Debugging is simplified by incremental testing because the modules and interfaces can be tested as the system grows.
- Programming errors in the interfaces and incorrect assumptions between units are discovered at an early stage.
- Because previously tested units are retested as new units are added, the probability of discovering less obvious errors is increased substantially.

4.4 Testing for Reentrancy

It is important to test your procedures for reentrancy before placing them into a library facility. Because ASTs can occur at any time, procedures that are not AST-reentrant may exhibit unexpected behavior. In particular, an AST occurring during storage modification in a procedure that is not AST-reentrant can corrupt the contents of the procedure's storage. (For further information about AST reentrancy, see Section 3.3.)

Full reentrancy is important to multi-thread tasking environments such as the environment used by Ada.

To avoid problems with reentrancy, carefully read and follow the coding guidelines described in Section 3.3.

4.4.1 Checking for AST Reentrancy

There are two methods of checking a procedure for AST reentrancy. You can use the Symbolic Debugger or perform a manual *desk check*.

4.4.1.1 Checking for AST Reentrancy using the Debugger

The following five steps are necessary when using the Debugger to check for AST reentrancy:

- 1 Create an activation of the procedure.
- 2 Set watchpoints on all storage used by the procedure.
- 3 Create a second activation of the procedure using the CALL command. Allow this second activation to run to completion. (The second activation represents the AST-level thread.)

Check to be sure that the AST-level thread of execution does not modify the storage accessed by the non-AST level thread of execution. If the AST-level thread of execution does modify any of that storage, check to ensure that it does not cause any unwanted side effects for the non-AST level thread of execution.
- 4 Step one instruction in the first activation.
- 5 Repeat steps 3 and 4 until the end of the procedure for the first activation.

For further information about the Debugger, refer to the *VMS Debugger Manual*.

4.4.1.2 Checking for AST Reentrancy by Desk Checking

Desk checking is the term for tracing through a procedure's execution manually. Performing a desk check for AST reentrancy consists of the following four steps:

- 1 Create an activation of the procedure being tested and its data using the method you normally use for manually tracing through a procedure.

This activation represents the non-AST level of your procedure's execution.
- 2 Create a second activation of the procedure using the process you used above. This second activation represents the AST-level thread of your procedure's activation.

Trace through the AST-level thread's execution to completion, one statement at a time.

Remember to update the contents of all storage locations and variables for each instruction of the procedure.

Check to be sure that the AST-level thread of execution does not modify the storage accessed by the non-AST level thread of execution. If the AST-level thread of execution does modify any of that storage, check to ensure that it does not cause any unwanted side effects for the non-AST level thread of execution.
- 3 Step through a single statement of the non-AST level thread of execution, remembering to update the contents of all storage locations.

Testing

4.4 Testing for Reentrancy

- 4 Repeat steps 2 and 3 until you have stepped through every statement in the non-AST level thread of execution. (Note that every statement of the AST-level thread is stepped through in each pass through step 2.)

As you can see, what you are actually doing in the process is testing between the execution of every two statements in the procedure. The most rigorous method of applying this type of desk checking for AST reentrancy is to step through the procedure at the assembly language level and test between each assembly language instruction.

4.4.2 Checking for Full Reentrancy

Full reentrancy differs from AST reentrancy in the number of threads of execution. An AST-reentrant environment can support only two threads of execution, the AST-level thread and the non-AST level thread. Full reentrancy is important in environments that can support many threads of execution, such as VAX Ada.

A procedure is fully-reentrant if any number of threads of execution can execute to completion without affecting any of the other threads of execution.

Generally, a procedure that is AST-reentrant is also fully reentrant. For further information on full reentrancy and environments supporting multiple threads of execution, refer to the documentation for VAX Ada.

4.5 Performance Analysis

All timer and resource allocation procedures should make statistics available for performance evaluation and debugging. You should code timer and resource allocation procedures with the following two entry points:

LIB_SHOW_name
LIB_STAT_name

4.5.1 SHOW Entry Point

A SHOW entry point provides formatted strings containing the information you need. The calling sequence for a SHOW entry point is as follows:

LIB_SHOW_name [code [,action-routine [,user-arg]]]

code

An optional code (of the form LIB_K_code) designating the statistic you need. Define a separate code for each statistic available; the codes should be the same for the SHOW and STAT entry points. The values associated with the codes start at one for each procedure. The functional specification in the procedure's documentation should list the codes used. If the code is omitted, or zero, the procedure provides all statistics.

action-routine

The address of an action routine. This is an optional argument. If omitted, statistics are written to SYS\$OUTPUT.

user-arg

An optional user argument to be passed to the action routine. If omitted, a shortened list is passed to the action routine. The **user-arg** argument, if present, is copied to the argument list passed to the action routine. That is, the 32-bit argument list entry passed by the calling program is copied to the argument list entry passed to the action routine. Thus, the access type, data type, argument form, and passing mechanism can be arbitrary, as agreed between the calling program and the action routine.

The optional action routine should have the form:

ACTION-ROUTINE (string [,user-arg])

See Section 3.1.6 for an example of the code to invoke a user action routine.

4.5.2 STAT Entry Point

A STAT procedure returns the information you want as binary results. The calling sequence is as follows:

LIB_STAT_name (code ,value)

code

A code designating the statistic you want. A separate code is defined for each statistic available; the codes are the same for the SHOW and STAT entry points. Codes start at one.

value

Is the value of the returned statistic.

4.6 Monitoring Procedures in the Run-Time Library

There are several procedures available in the Run-Time Library for time and resource monitoring. These Run-Time Library procedures and their functions are as follows:

- **LIB\$SHOW_VM**

LIB\$SHOW_VM is a resource monitoring procedure that returns the statistics accumulated from calls to LIB\$GET_VM and LIB\$FREE_VM.

The following three statistics are returned by default:

- Number of successful calls to LIB\$GET_VM
- Number of successful calls to LIB\$FREE_VM
- Number of bytes allocated by LIB\$GET_VM but not yet deallocated by LIB\$FREE_VM

LIB\$SHOW_VM returns these statistics in the formatted form, nnnn.

- **LIB\$STAT_VM**

LIB\$STAT_VM is a resource monitoring procedure that returns to its caller one of the three statistics available from calls to LIB\$GET_VM and LIB\$FREE_VM. These are the same statistics that are returned by LIB\$SHOW_VM. Unlike LIB\$SHOW_VM, which returns the statistics in formatted form to SYS\$OUTPUT, LIB\$STAT_VM returns the specified statistic in a signed longword integer.

Testing

4.6 Monitoring Procedures in the Run-Time Library

- **LIB\$SHOW_TIMER**

LIB\$SHOW_TIMER is a time monitoring procedure that returns the times and counts accumulated since the last call to LIB\$INIT_TIMER and displays them on SYS\$OUTPUT. A user-supplied action routine may alter this default behavior.

The following five statistics are provided by default:

- Elapsed real time
- Elapsed CPU time
- Count of buffered I/O operations
- Count of direct I/O operations
- Count of page faults

- **LIB\$STAT_TIMER**

LIB\$STAT_TIMER is a time monitoring procedure that returns the same information as LIB\$SHOW_TIMER. The difference is that LIB\$STAT_TIMER returns the information as an unsigned longword or quadword, whereas LIB\$SHOW_TIMER returns the information in the format hhhh:mm:ss:cc for times and the format nnnn for counts. In addition, LIB\$STAT_TIMER returns only one of the five available statistics per call.

For further information about these time and resource monitoring procedures, see the *VMS RTL Library (LIB\$) Manual*.

5

Integration

5.1 Grouping Procedures

Modular procedure libraries consist of compiled and assembled object code intended to be associated with a calling program at link time. References to procedures in these libraries are resolved when the linker searches the user libraries specified in the LINK command or when it searches the default system libraries. The program can then call library procedures at run time.

DIGITAL supplies several procedure libraries, such as the VAX Common Run-Time Procedure Library (also called the Run-Time Library), that support components of the VMS operating system. You can explicitly access the procedures in the Run-Time Library to perform frequently used operations. To do this, simply include calls to Run-Time Library procedures in your program. The linker automatically searches the default libraries to resolve references to Run-Time Library procedures. (For information on what procedures are available in the Run-Time Library, see Section 1.3.1.)

You can create your own procedure libraries and shareable images by following the guidelines in Section 5.1.2 and Section 5.2.1. Section 5.2.2 shows how to use transfer vectors to make your shareable images easier to maintain. Before you begin grouping modular procedures, make sure they conform to the rules listed in Appendix A.

There are three ways to group modular procedures:

- Combine object modules into an object module library.
- Link object modules together into a shareable image.
- Combine shareable images into a shareable image library.

The following sections show how to create and install these three types of procedure libraries and how to access them when you link and run a program.

5.1.1 Creating Facility Prefixes

The facility prefix is the group identifier for a set of related procedures contained in a library facility. The facility prefix appears in the procedure name of every procedure in that library facility. An example of a library facility is the Screen Management Facility in the Run-Time Library. The names of all procedures appearing in the Screen Management Facility begin SMG; for example, SMG\$ERASE_CHARS.

To create your own facility prefix, follow these steps:

- 1 Choose a facility prefix. This prefix can be from 1 to 27 characters in length. However, it is recommended that you choose facility prefixes between 2 and 4 characters.

Integration

5.1 Grouping Procedures

- 2 If your facility will be generating messages, you must specify a unique facility number in the message source file. This number can range from 0 to 4095. Any number within this range and not being used by someone else on your system is acceptable. This facility number will be used by the message utility in generating the condition value for the message.

Bit 27 (STS\$V_CUST_DEF) of a condition value indicates whether that value is supplied by the user or by DIGITAL. This bit must be 1 if the facility number is user created. For further information, see the *VMS System Messages and Recovery Procedures Reference Volume*.

- 3 Use the facility prefix when naming all procedures within the new facility. Remember to follow the naming conventions described in Section 3.1.2.

5.1.2 Creating Object Module Libraries

In addition to using the system default object module libraries, you can also create your own object module libraries. An object module library created by you can contain object files produced by any VAX language compiler.

The following three steps are required to create an object module library:

- 1 Write or collect the procedures you want to group together in the object module library. Make sure the names chosen for the modules, procedures, and facility prefixes conform to the guidelines given in Section 3.1.2.
- 2 Compile the source code to create the object files for the procedures to be contained in the library.
- 3 Create the object module library using the LIBRARY command.

The format of the library command is as follows:

```
$ LIBRARY /CREATE library-name.OLB filespec1.OBJ -  
_$ [,filespec2.OBJ [...]]
```

The parameter **library-name** is the name that you have given the library. The default file extension for **library-name** is OLB.

The **filespec** parameters are the object files for the procedures you want the object module library to contain. The default file extension for **filespec** is OBJ.

To clarify this process, create a sample object module library.

The first step is to choose the procedures to be contained in the sample library. Call the sample library GRAPHICS.OLB. This library will contain modular procedures for creating mathematical representations of circles, cylinders, squares, and other geometric shapes.

The files to be used are GRASPHERE.BAS and GRACUBE.FOR. GRACUBE.FOR contains a single procedure, GRA_CUBE, to generate cube shapes. (Note that GRA is the facility prefix, and the underscore in the procedure name indicates that this is a user-defined procedure.) GRASPHERE.BAS contains several procedures that are grouped together because they share similar code. The procedures contained in GRASPHERE create spheres (GRA_SPHERE), oblate spheroids (GRA_OBL_SPH), and spherical sections (GRA_SPH_SEC).

Now that you have chosen the procedures and named them correctly, advance to step 2 and create the object files necessary to build the object module library. The commands used to create the object files are as follows:

```
$ BASIC GRASPERE.BAS  
$ FORTRAN GRACUBE.FOR
```

This produces the object files GRASHPERE.OBJ and GRACUBE.OBJ.

The final step is to create the object module library itself from those object files. Create the object module library by entering the following command:

```
$ LIBRARY /CREATE GRAPHICS GRASPERE,GRACUBE
```

Note that you did not specify the file extensions because you used only the default values.

Once the LIBRARY command has been entered, the object module library GRAPHICS.OLB is ready to be linked with an application program.

Figure 5-1 shows the overall development of the user-created library of graphics procedures, GRAPHICS.OLB.

5.2 Shareable Library Images

If you have a collection of procedures you expect a number of users to use, you may want to group these procedures into a shareable library image. A shareable library image, usually referred to as a shareable image, is similar to an object library except that it has been pre-linked so that all references between procedures in the library have already been resolved.

A shareable library image has the following advantages:

- Conserves memory space

Several processes can "share" a single copy of a shareable image rather than each process retrieving its own copy from the disk.

- Conserves disk storage space

Programs linked to a shareable library image share a single disk copy of the library code rather than each program including the code in its own executable image.

- Shortens link time

Since the internal references in the library have already been resolved, there is less work for the linker.

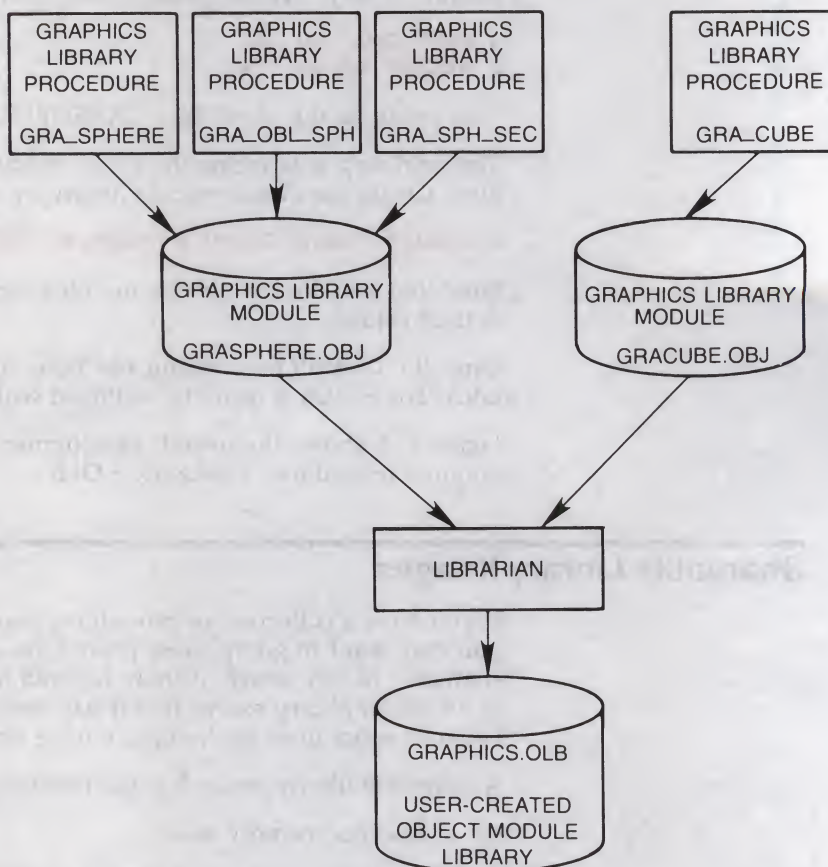
- Allows for updates without relinking

You can supply a new version of a shareable library image that can automatically be used by all programs linked to it without the need for the users to relink their programs.

Integration

5.2 Shareable Library Images

Figure 5-1 Development of a User-Created Object Module Library



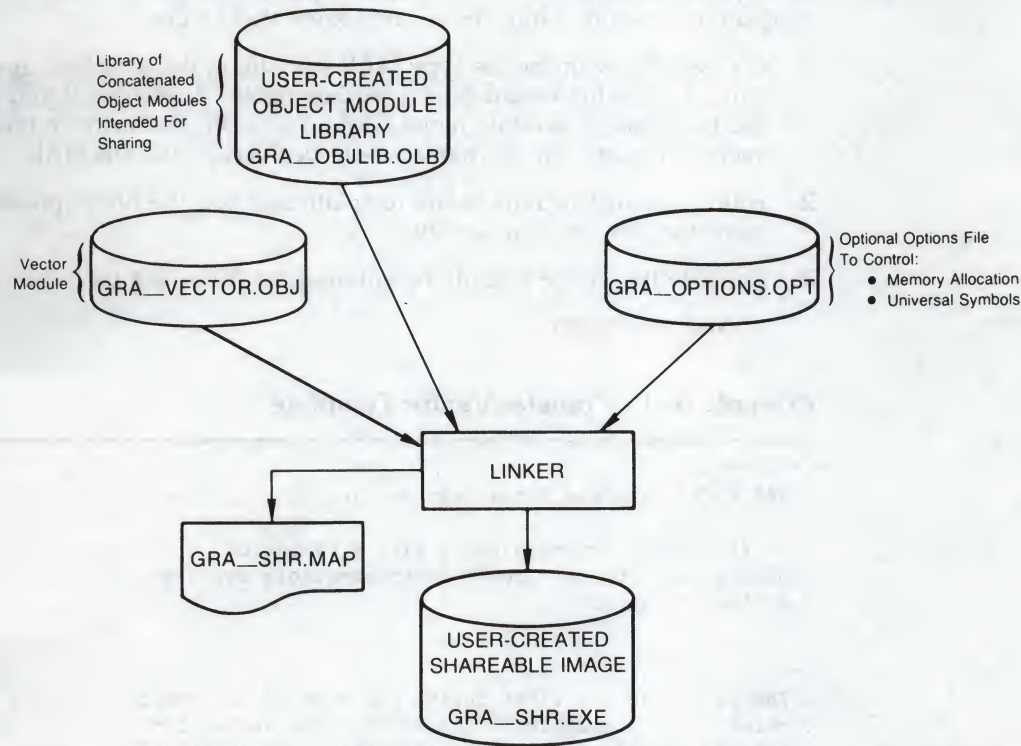
ZK-4028-85

The greatest benefit of using a shareable library image is the ability to conserve physical memory by sharing a single memory copy of the library code among all users. This feature is enabled by having a privileged user, usually the system manager, use the VMS Install Utility to install the shareable library image with the /SHAREABLE attribute. This benefit is not limited to shareable library images, however. Any executable image can be installed as shareable. See the *VMS Install Utility Manual* for more information.

5.2.1 Creating Shareable Library Images

It is quite simple to take a collection of procedures and link them together as a shareable library image. In addition to your procedures, you need to provide a transfer vector and a linker options file. The following sections show how to create both of these items in a step-by-step fashion. If you follow the "recipe", you will be rewarded with a shareable image library that is easy to use and maintain.

Figure 5-2 Creating a Shareable Image



ZK-4029-85

Figure 5-2 provides an overview of the process of creating a shareable image.

5.2.2 Creating the Transfer Vector

When a program links to a shareable library image, the linker stores the following items of information in the program about each reference to a routine in that image:

- The name of the shareable image
- The location of the routine entry point relative to the beginning of the shareable image

Since a relative location is used, the routine cannot change its location without making all images using this routine invalid. Using a transfer vector allows you to modify code and rearrange routines without invalidating previously linked programs.

A transfer vector, which is placed at the beginning of the image, is a module containing a list of "forwarding addresses" for routines in the shareable image. When a transfer vector is used, programs linking to the shareable image reference the entry points in the transfer vector rather than the actual routines. The entry in the transfer vector then transfers control to the actual routine. Because each entry in the transfer vector is the same size, it is easy to keep the relative location of an entry the same across updates.

Integration

5.2 Shareable Library Images

A transfer vector must be created using the VAX MACRO assembler language, but it is not necessary to know VAX MACRO to create a transfer vector.

A template for a transfer vector is provided in Example 5-1. To use this template in creating a transfer vector, follow these steps:

- 1 Create a file with the file type MAR containing the template. A suggested name for the file would be the module name. Therefore, if you were to use the example module name GRA_VECTOR, as shown in the transfer vector template, the file name would be GRA_VECTOR.MAR.
- 2 Follow the instructions in the template and edit the file to produce the correct vector for your facility.
- 3 Compile the vector module by entering the following DCL command:

```
$ MACRO GRA_VECTOR
```

Example 5-1 Transfer Vector Template

```
;+
; VAX MACRO template for a transfer vector.
;
; In VAX MACRO, comments begin with a semicolon.
; Blanks and tabs may appear interchangeably wherever
; a blank is shown.
;-

;+
; The following two lines define the name of the module,
; which in this example is GRA_VECTOR (the vector for
; the GRA_ facility), and the identification (version) of
; this module, which is 1-001. Replace GRA_VECTOR by
; whatever module name you choose, and replace 1-001
; by whatever version number you choose.
;-

        .TITLE  GRA_VECTOR
        .IDENT  /1-001/

;+
; The following lines define a MACRO called ROUTINE which,
; when invoked, generates one transfer vector entry. Enter
; these lines exactly as shown; no customization is necessary.
;-

        .MACRO  ROUTINE NAME
        .EXTRN          NAME
        .ALIGN  QUAD
        .TRANSFER      NAME
        .MASK           NAME
        JMP             NAME+2
        .ENDM

;+
; The following lines define a MACRO called FUTURE which,
; when invoked, generates one transfer vector entry that does
; not have a universal symbol in the shareable library image.
; It's used to "reserve" transfer vector entries for routines
; that haven't been written yet, so the length of the transfer
; vector doesn't have to change when they're written.
;-
```

Example 5-1 Cont'd. on next page

Example 5-1 (Cont.) Transfer Vector Template

```

.MACRO FUTURE NAME
.EXTRN      NAME
.ALIGN QUAD
.MASK      NAME
JMP        NAME+2
.ENDM

;+
; The next two lines define the program section (PSECT)
; in which the transfer vector resides. Change the name
; of the PSECT, GRA_VECTOR, to whatever program section
; name you choose. The remainder of the two lines define
; PSECT attributes which you should not change.
;
; When linking the shareable image that uses this transfer
; vector, it's important to place the transfer vector at
; the beginning of the shareable image (for example, by
; means of the CLUSTER or COLLECT options in a linker options
; file. See example 5-4.
;-
.PSECT GRA_VECTOR PIC, USR, CON, REL, LCL, SHR, -
                    EXE, RD, NOWRT, QUAD

;+
; Following this point are the transfer vector entries.
; Each entry is of the form:
;
ROUTINE routine-name
;
; where routine-name is the name of the routine for which you
; wish to make an entry. The order in which you list the routines
; is not important, but it is important that you not change the
; order once you have created the shareable library image.
;
; Note that there are 'not implemented' entries at the end of
; the vector. These are placeholders in the vector for routines
; to be added in the future. As long as the routines do not exist,
; these placeholders should transfer to a routine that signals, or
; returns an error. When you add routines, simply replace a FUTURE
; entry with a ROUTINE entry with the name of a new routine.
;
; If you run out of 'not implemented' entries, then extend the
; vector with a new set, and increment the minor ID in the GSMATCH
; option in the linker options file (see example 5-4).
;
; For compatibility, you should never delete a routine. If the
; routine must be eliminated, replace its name with the name of
; the 'not implemented' routine, so that an error will be signalled
; if a program (mistakenly) tries to use it.
;
; As examples, entry declarations for the routines GRA_SPHERE,
; GRA_OBL_SPH, GRA_SPH_SEC and GRA_NOT_IMPLEMENTED are shown.
; Replace these by your own entries.
;-

```

Example 5-1 Cont'd. on next page

Integration

5.2 Shareable Library Images

Example 5-1 (Cont.) Transfer Vector Template

```
ROUTINE GRA_SPHERE
ROUTINE GRA_OBL_SPH
ROUTINE GRA_SPH_SEC
FUTURE  GRA_NOT_IMPLEMENTED
FUTURE  GRA_NOT_IMPLEMENTED
```

```
;+
; The last line denotes the end of the module.
;-

.END
```

5.2.3 Creating the Linker Options File

A linker options file contains instructions to the linker about how it should build the executable or shareable image. The use of a linker options file is rarely necessary for executable images because the linker's defaults are usually adequate. However, shareable images need additional information, which can only be supplied by an options file.

A linker options file can specify many different image attributes, but only a few of these are usually needed for shareable library images. Example 5-2 shows a linker options file template and briefly describes each option that is used. To customize this linker options file for your own needs, follow the instructions in the options file comments. A suggestion for the name of the file would be the facility prefix followed by `_OPTIONS`, with a file type of `OPT`. Thus, for the example GRA facility, the options file name would be `GRA_OPTIONS.OPT`.

Example 5-2 provides a template for a linker options file.

Example 5-2 Template for a Linker Options File

```
!+
! In a linker options file, comments begin with an exclamation
! point. Blanks and tabs may appear interchangeably wherever
! a blank is shown.
!-

!+
! The first line defines the name of your shareable library image.
! In this example the name is GRA_SHR. Replace GRA_SHR with the
! name of your shareable library image.
!
! If you don't specify the image name, the Linker will use the
! name of the first input file.
!-

NAME = GRA_SHR
```

Example 5-2 Cont'd. on next page

Integration

5.2 Shareable Library Images

Example 5-2 (Cont.) Template for a Linker Options File

```
!+
! The next line specifies the identification string. The
! identification string is the version number of the shareable
! library image. Replace "V1.0" with the identification string of
! your shareable image.
!
! (You can display this identification string later by using
! the ANALYZE/IMAGE command.)
!-

IDENTIFICATION = "V1.0"

!+
! The following line specifies the global section match control.
! Match control can be used to control whether programs previously
! linked with this shareable library image need to be relinked.
! The control consists of three parts; match criteria, a major ID
! and a minor ID.
!
! The major ID can be changed to force a program to be relinked.
! In this example, the major ID is 1. Unless you need to force
! programs linked with this shareable library image to be
! relinked, do not change the major ID.
!
! The minor ID is used in combination with the match criteria to
! determine whether a new link operation is necessary as well.
! In this example, the match criteria is "LEQUAL" (less than or
! equal). This means that any program linked with a version of
! this shareable library image that has a minor ID less than, or
! equal to this minor ID, may run. Programs linked with a version
! of this shareable library image with a minor ID greater than
! this minor ID may not run. The minor ID should be changed when
! the size of the transfer vector is changed (increased). You
! may wish to change the minor ID value with each version of the
! shareable library image (whether the size of the transfer vector
! has changed or not), but that is not necessary.
!
! See the VMS Linker Utility Manual for more information.
!-

GSMATCH = LEQUAL,1,1

!+
! If you have blocks of data defined with COMMON (FORTRAN, BASIC,
! PASCAL), MAP (BASIC), GLOBALDEF (PL/I), PSECT_OBJECT (Ada), or
! similar declarations in other languages, you may wish to change
! certain attributes of the PSECTs (Program Sections) defined by
! these language features.
!
! In this example, we assume that there is a COMMON named GRA_COMMON,
! and that we wish to change it's SHR attribute to NOSHR. We also
! want to modify the PSECT attributes of read-only data generated by
! FORTRAN so that it can be collected into the same image cluster
! with the code (read-only data may be shared).
!-

PSECT_ATTR = GRA_COMMON,NOSHR
PSECT_ATTR = $PDATA,PIC,EXE
```

Example 5-2 Cont'd. on next page

Integration

5.2 Shareable Library Images

Example 5-2 (Cont.) Template for a Linker Options File

```
!+
! The next three lines group the PSECTs contributing to the
! shareable library image into clusters. PSECTs with similar
! attributes may be collected into the same cluster. Note that
! the transfer vector is always placed so that it comes at the
! beginning of the shareable library image.
!
! In this example, the transfer vector resides in its own cluster.
! Code and read-only data occupy a second one, and the remaining
! PSECTs make up the third cluster.
!
! Replace GRA_VECTOR by the PSECT name you specified for your
! transfer vector. The other PSECT names shown here are typical
! of those generated by high-level languages. You should check
! your compilation listings to determine PSECTs used, their
! attributes, and group them accordingly.
!-

COLLECT = CLUSTER1, GRA_VECTOR
COLLECT = CLUSTER2, $CODE, $PDATA
COLLECT = CLUSTER3, $LOCAL, $BLANK, GRA_COMMON

!+
! The following lines list the modules that are to be
! included in the shareable library image. This
! example assumes that all of the necessary object
! modules are in an object module library called
! GRA_OBJLIB.OLB.
!
! If you need to reference more modules or libraries, add new
! lines containing the names of the appropriate files.
!-

GRA_OBJLIB.OLB/LIBRARY

!+
! End of linker options file.
!-
```

For more information about linker options files and the linking process, see the *VMS Linker Utility Manual*.

5.2.4 Creating the Shareable Library Image

Once you have created the transfer vector and the linker options file, you are ready to link the shareable library image. Using the example of a shareable library image for the GRA facility, GRA_SHR, enter the following command to link the image:

```
$ LINK/SHAREABLE=GRA_SHR.EXE/MAP=GRA_SHR.MAP/FULL GRA_OPTIONS/OPTIONS
```

Substitute your own image name and linker options file name for the names GRA_SHR and GRA_OPTIONS above. The linker creates the shareable library image as GRA_SHR.EXE and also creates a link map file, containing useful information about the linked image in a file named GRA_SHR.MAP. If you also want a cross-reference of symbols used in your shareable library image, add the command qualifier /CROSS_REFERENCE to the command line.

The link should complete without any error or other messages being displayed. There is one warning that may occur if any of the modules in your library are written in either VAX MACRO or VAX BLISS-32. (Programs written in high-level languages do not exhibit this problem.) The text of the message is "basing image due to errors in relocatable references." While the link operation completes successfully and the resulting shareable library image can be used, you lose the desirable attribute of upwards compatibility when you link a new image or when an image referenced by your shareable library image is relinked. The cause of this error is references to external symbols without using general mode addressing. In VAX MACRO, prefix all external references with G`. In VAX BLISS, either declare the external name with the attribute ADDRESSING_MODE (general) or use the SWITCHES statement to force general-mode addressing for externals. See the *VAX MACRO and Instruction Set Reference Manual* and the *BLISS Language Guide* for more information.

Once you have linked your image, examine the linker map to make sure that your transfer vector (in this example PSECT \$\$\$\$VECTOR) has been allocated at a base address of zero. If it has not, you will probably need to change the attributes of some PSECTs to PIC and EXE, as shown in the template linker options file. See Section 5.3.4.1, Generation of Image Sections, in the *VMS Linker Utility Manual* for additional information.

5.2.5 Combining Shareable Images into a Shareable Image Library

Given a collection of shareable images, you can create a shareable image library using the library command. The format of the library command used to create a shareable image library is as follows:

```
$ LIBRARY /CREATE /SHARE library-name -
_$ shareable-image1 [,shareable-image2 [...]]
```

The parameter **library-name** is the name that you have given the shareable image library. The **shareable-image** parameters are the shareable images you want the shareable image library to contain.

5.3 Linking to Libraries of Modular Procedures

To use a library of modular procedures, whether in an object module library or in a shareable image library, specify the name of the library along with the /LIBRARY file qualifier as one of the input files on the LINK command. For example:

```
$ LINK MYPROG,OURLIB/LIBRARY
```

In this example, MYPROG.OBJ is the main program of what will be an executable image, and OURLIB.OLB is either an object module library or a shareable image library. However, if OURLIB is a shareable image library and if the shareable library image does not reside in the system directory SYS\$LIBRARY, then you must define a logical name for the shareable image name to point to the file's actual location, as in the following example:

```
$ DEFINE GRA_SHR OURDISK:[OURLIB]GRA_SHR
```

This logical name must be defined before the program is run.

The teacher's role is a complex one, involving a wide range of responsibilities. In the classroom, the teacher is responsible for creating a supportive learning environment, delivering instruction, and assessing student progress. Beyond the classroom, the teacher may also be involved in curriculum development, professional development, and community outreach. The teacher's role is essential to the success of any educational institution.

The teacher's role is a complex one, involving a wide range of responsibilities. In the classroom, the teacher is responsible for creating a supportive learning environment, delivering instruction, and assessing student progress. Beyond the classroom, the teacher may also be involved in curriculum development, professional development, and community outreach. The teacher's role is essential to the success of any educational institution.

2.2. Combining Spreadsheets into a Separate Image Library

Combining spreadsheets into a separate image library is a useful technique for organizing data. This process involves creating a central repository for all spreadsheet data, which can then be accessed and analyzed as needed. This approach offers several advantages, including improved data management, enhanced security, and increased flexibility in data analysis.

3.3. Linking to Libraries of Molecular Structures

Linking to libraries of molecular structures is a critical step in the process of molecular modeling. This involves establishing a connection between a specific molecular structure and a library of related structures. This approach allows researchers to quickly access and analyze a wide range of molecular data, facilitating the discovery of new insights into molecular behavior and interactions.

6 Maintenance

6.1 Upward Compatibility

Upward compatibility is a very important concept in the maintenance stage. If a procedure is upwardly compatible, then changes and updates to the procedure do not affect the execution and use of the previous versions of that procedure.

For example, imagine a user-written procedure named LIB_TOTAL_BILL. The calling sequence for this procedure is as follows:

```
CALL LIB_TOTAL_BILL (sale, tax)
```

Assume that the user who wrote this procedure decided to update the procedure so that it could be used to calculate the total bill for credit card customers. To do this, a third argument, **interest**, must be added. To be upwardly compatible, adding the argument **interest** must not conflict with the way the procedure was previously run. The new calling sequence would be as follows:

```
CALL LIB_TOTAL_BILL (sale, tax [,interest])
```

The procedure should be written so that the user can still call the procedure as it was called before, simply omitting the **interest** argument.

If, in the updated version of this procedure, the user can still follow the calling sequence of the previous versions, the procedure is said to be upwardly compatible.

6.1.1 Making Your Procedures Upwardly Compatible

To be compatible with all future versions of the shareable image, shareable image procedures must adhere to the following rules (in addition to following the VMS Modular Programming Standard):

- A procedure's entry point is referenced through a transfer vector. (For further information, see Section 5.2.2.)
- A procedure's code and data are position independent.

6.1.2 Regression Testing

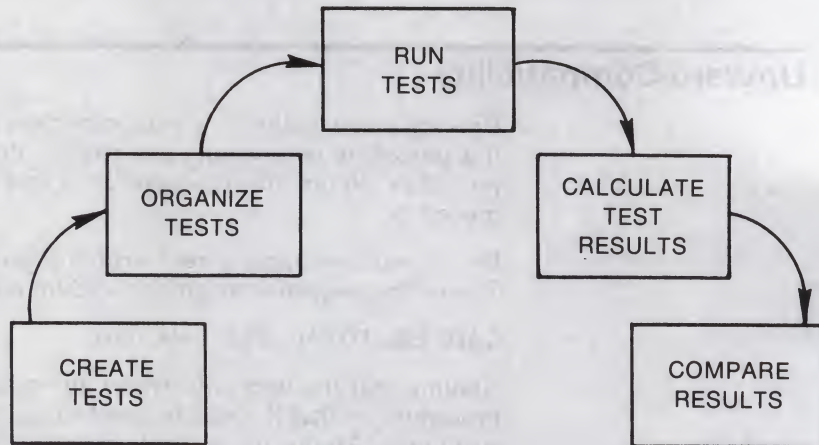
Regression testing is a method of ensuring that new features added to a procedure do not affect the correct execution of previously tested features. In regression testing, you run established software tests and compare test results with expected results. If the actual results do not agree with what you expected, the software being tested may have errors. If errors do exist, the software being tested is said to have "regressed."

Maintenance

6.1 Upward Compatibility

The following steps, as illustrated in Figure 6-1, are involved in regression testing:

Figure 6-1 Regression Testing



ZK-4061-85

1 Create tests

Create tests by writing command files to test your software.

2 Organize tests

Organize test files to allow easy access to tests as they are needed.

3 Run tests

To run a single test, submit its command file to the batch queue.

To run multiple tests, create a command file that submits each test to the batch queue.

4 Calculate test results

Calculate the expected test results either by hand or by using previously tested software.

5 Compare results

Compare the actual test results to the results you expected. If there are inconsistencies, repeat your calculation in step 4. If the inconsistency still exists, examine the changes you have made to the software to discover the error.

It is important to write new tests and repeat the regression testing steps every time you add new functionality to the procedure. If you do not do so, the procedure may regress while the errors go undetected.

6.1.2.1 Updating the Transfer Vector

The following three steps are required to update the transfer vector:

1 Add the new entry or entries.

To maintain upward compatibility when updating a procedure library, it is important not to disturb the order of existing entries in the transfer vector. Always add new entries to the end of the transfer vector and do not delete or rearrange existing entries. If you carefully follow these guidelines (which are also given in Section 5.2.2), you should have no problems updating the transfer vector, and your shareable image will be upwardly compatible.

2 Check the transfer vector if you suspect that there is a compatibility problem.

To check the transfer vector, compare the entry addresses found in the listing file generated by compiling the transfer vector with the entry addresses in a listing file from the previous version of the transfer vector. Therefore, always keep a copy of previous listing files. In comparing files, make sure that the addresses of the symbols have not changed.

If all symbols in the previous version are contained in the current version and the relative addresses of those symbols have not changed, then the transfer vector has been changed in an upwardly compatible fashion and existing programs using symbols within the transfer vector continue to be valid.

3 Update the minor global section match identification in the linker options file.

See the linker options file template in Section 5.2.3.

Note that if you pre-extend the transfer vector with dummy entries, then GSMATCH does not have to be incremented.

6.2 Adding Arguments to Existing Routines

During the normal course of maintenance, it sometimes becomes necessary to pass new or additional information to an existing procedure rather than create a new procedure. This new information may be passed to the procedure in one of the following two ways:

- Directly, by adding new arguments to the procedure
- Using an argument block

Maintenance

6.2 Adding Arguments to Existing Routines

6.2.1 Adding New Arguments to the Procedure

There are two rules you must follow when directly adding new arguments to a procedure:

- New arguments must be added at the end of the existing argument list.
- New arguments must be optional.

It is important that new arguments be added at the end of the existing argument list to maintain upward compatibility. If you change the order of the existing arguments by placing the new argument at the beginning or middle of the list, all applications written with the previous version of the procedure will no longer work.

Your procedure should also treat the new argument as an optional argument. If the new argument is required, applications that used the previous version of the procedure are invalidated.

Because you cannot assume that all previously written applications will be rewritten to include the procedure's new argument, the procedure must test for the argument's presence before attempting to access it. If the procedure does not verify the presence of the new argument and attempts to access that argument when it is not present, the results will be unpredictable.

The passing mechanism of the new argument must conform to the guidelines established in Section 2.2.1.

6.2.2 Using Argument Blocks

By using an argument block, you can avoid adding multiple arguments to your procedure. When an argument block is used, the calling program passes a single argument to the called procedure. This argument is the address of an argument block. The argument block is a block of information containing any information agreed on by the calling and called procedures. This information is required by the called procedure in order to perform its task.

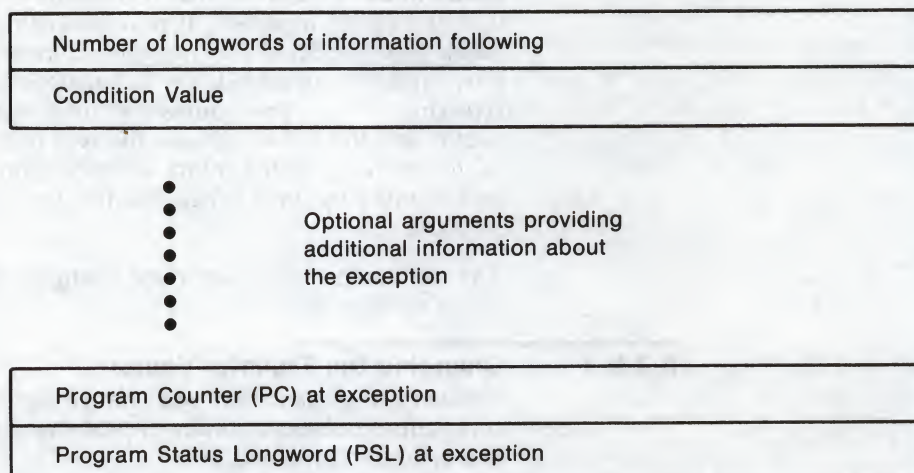
The argument block itself is simply a contiguous piece of virtual memory. The information contained in the argument block can be numeric or scalar data, descriptors, bit vectors, and so on. The format is simply agreed on by the users of the procedure and its writer.

The first longword in the argument block contains the length of the block. The length can be in bytes, longwords, or whatever, but it must be agreed on by both the calling program and the called process and be implemented and documented as such.

One example of an argument block is the signal argument vector used in condition handling. A condition handler is called with a signal argument vector and a mechanism argument vector. Each vector is an example of an argument block. The signal argument vector in Figure 6-2 is an example of an argument block.

6.2 Adding Arguments to Existing Routines

Figure 6-2 One Type of Argument Block, the Signal Argument Vector



ZK-4030-85

As you can see, the signal argument vector contains the number of longwords of actual information in its first longword. What information actually follows depends on the condition value of the signal.

Note that, if you lengthen an argument block to provide new information to a called procedure, your procedure should check the length of the argument block for validity before attempting to access the information. As with adding new arguments directly to a procedure, the calling program may have been written to pass the previous, shorter argument block. If your procedure does not check and attempts to access information past the end of the actual argument block, the results will be unpredictable.

6.3 Updating Libraries

Any time modifications or enhancements are made to modular procedures that are a part of some library, the library containing the procedures must be updated to reflect the new or changed procedures.

6.3.1 Updating Object Libraries

If the updated procedures are in an object library, the library needs to be updated so that subsequent access to that library by LINK or other commands will access the object modules for the new or changed procedures.

To update an object library, use the LIBRARY command with the REPLACE and OBJECT qualifiers, as follows:

```
$ LIBRARY /REPLACE library-name filespec[,...]
```

In this example **library-name** is the name you have given the library. The default file type for library-name is OLB. The name of an object module is **filespec**. The default file type for **filespec** is OBJ.

Maintenance

6.3 Updating Libraries

6.3.2 Updating Shareable Images

If the updated procedures are part of a shareable image, the shareable image needs to be relinked so that it contains the new or changed versions of any updated object modules. If new procedures are added, the transfer vector needs to be updated and recompiled prior to relinking the shareable image. If new modules are added, the linker options file needs to be updated prior to relinking. If new procedures and new modules are added, then the transfer vector and the linker options file will need to be updated. If the transfer vector is changed, the minor identification value of the GS MATCH must be incremented by one. When this has been done, the shareable image may be relinked.

The following sections describe changing the transfer vector and updating the linker options file.

6.3.2.1 Changing the Transfer Vector

It is very important that you change the transfer vector in an upwardly compatible fashion in order to preserve the validity of existing programs that use symbols contained in it.

To avoid invalidating the transfer vector, adhere to the following rules:

- Do not change the order of existing symbols in the transfer vector.
- Do not remove any existing symbols from the transfer vector.
- Add all new symbols to the end of the existing transfer vector.

Recompile the vector module when the changes have been completed, and examine the names and relative addresses assigned to the symbols in the module to ensure that all symbols previously there are still there and that all relative addresses of those symbols have not changed. If existing symbols are removed or relocated, then existing programs using the shareable image will be invalidated.

6.3.2.2 Updating the Linker Options File

Several items in the linker options file should be changed each time the transfer vector is updated.

- Increment the minor identification portion of the global section match identification string.

Changing the minor identification value prevents programs linked against the new version of the shareable image from activating an older version (which may not contain new procedures). Global section match identification values are set with the GSMATCH option.

- If necessary, add any new object modules to be included in the shareable image to the linker options file.
- Change the identification string to a new version number.

6.3.3 Updating Shareable Image Libraries

If the updated procedures are in a shareable image that is part of a shareable image library, the shareable image library needs to be updated so that subsequent access to that library by the LINK command will be to the updated version of the shareable image.

To update a shareable image library, use the LIBRARY command with the REPLACE and SHARE qualifiers, as follows:

```
$ LIBRARY /REPLACE /SHARE library-name filespec[,...]
```

In this example, **library-name** is the name you have given the library. The default file type for library-name is OLB. The name of a shareable image is **filespec**. The default file type for **filespec** is EXE.

Updating (Revisable) Inside Thinking

When you are in a negative state of mind, you are likely to have a negative outlook. This is because your mind is focused on the negative aspects of your life. To change this, you need to update your thinking. This means that you need to look at your life from a different perspective. You need to see the positive aspects of your life, even if they are small. You need to see the things that you are grateful for. You need to see the things that you are proud of. You need to see the things that you are capable of. You need to see the things that you are working on. You need to see the things that you are achieving. You need to see the things that you are becoming. You need to see the things that you are creating. You need to see the things that you are doing. You need to see the things that you are living. You need to see the things that you are being. You need to see the things that you are having. You need to see the things that you are doing. You need to see the things that you are living. You need to see the things that you are being. You need to see the things that you are having.

A VMS Modular Programming Standard

This appendix presents the VMS standard for writing modular procedures in any VAX language. The elements of this standard are the minimum necessary to interface your software at the callable procedure level with software written by others, and vice versa.

Nonconformance to any elements in this standard must be indicated in your procedure's documentation.

Each element of the standard is described in greater detail in other sections of this manual. References to the appropriate sections appear after each element. The word "Optional" appears before the section reference if adherence to the element is not required to maintain modularity.

A.1 Purpose of this Standard

Procedures can be combined to form programs in the following ways:

- Your procedure calls other procedures
- Other procedures call your procedure
- A calling program calls either your procedure or other procedures

In order for procedures to execute as expected when combined to form a program, general agreements among programmers are necessary. Modular procedures that do not follow these general agreements could cause other procedures in the program image to execute incorrectly.

This VMS Modular Programming Standard is designed to give programmers a common environment in which to write their code. If all programmers follow this standard, then any modular procedure can be added to a procedure library without conflicting with any procedures already in the library or with any that might be added later.

A.2 Applicability

The elements of this standard apply to library procedures and are recommended for other types of software, including utilities and application programs. Each DIGITAL-supplied programming language implemented on the VMS operating system lets you write your procedures to follow this standard.

This standard applies to procedures that have a public entry point. A public entry point is one that the linker is able to locate by searching the default system libraries. This standard does not apply to calls to routines internal to a module that do not have public entry points, as long as the entire set of procedures follows the standard.

VMS Modular Programming Standard

A.3 Coding Rules

A.3 Coding Rules

The coding rules pertain to all procedures. These rules are grouped in the following categories:

- The Calling Interface
- Initialization
- Reporting Exception Conditions
- AST Reentrancy
- Resource Allocation
- The Format and Content of Coded Modules
- Shareable Images
- Upward Compatibility

Detailed descriptions of the rules for each of these categories are presented in the sections that follow.

A.3.1 The Calling Interface

- Calls to procedures must follow the VAX Procedure Calling and Condition Handling Standard. Some elements of this standard restrict procedures to a subset of the VAX Procedure Calling and Condition Handling Standard to increase the ability of procedures to call each other. (See *Introduction to VMS System Routines*.)
- A procedure makes no assumptions about its environment other than those of this standard. In particular, to operate as specified, a procedure neither makes assumptions about, nor places requirements on, the calling program.
- A procedure should not call other procedures or system services if the resulting combination violates this standard from the calling program's viewpoint. A procedure can call other procedures or system services that do not follow optional elements of this standard. However, if the resulting combination (as seen from the calling program) does not follow the optional elements, the calling procedure must indicate such nonconformance in its documentation. (See Section 3.1.5.)
- A modular procedure must provide an interface to its callers that allows the callers to follow all required elements of this standard.
- Each module should contain only a single public entry point. (Optional.)
- When a procedure uses a JSB entry point, it should also provide an equivalent call entry point to maintain language-independence. This is because, although JSB calling sequences may execute faster than procedure calls, an explicit JSB linkage to an external routine may not be provided in some high level languages. (Optional. See Section 2.3.)
- The order of required arguments should be the same as that of the VAX hardware instructions, namely, read, modify, and write. Optional arguments follow in the same order. However, (according to the VAX Procedure Calling and Condition Handling Standard) if a function value cannot be represented in 64 bits or is of type string, the first argument

VMS Modular Programming Standard

A.3 Coding Rules

specifies where to store the function value, and all other arguments are shifted one position to the right. (See Section 2.2.4.)

- A procedure's caller should indicate omitted trailing optional arguments either by passing argument list entries that contain zero or by passing a shortened argument list. However, system services require trailing arguments and do not adhere to this guideline. (Optional. See Section 2.2.5.)
- String arguments should always be passed by descriptor. (See Section 4.2.)
- Procedures must not accept data from, nor return data to, their calling programs by using implicit overlaid PSECTs or implicit global data areas. All arguments accepted from or returned to the calling program must use the argument list and function value registers (R0 and R0/R1). (See Section 2.2.2.)
- A procedure cannot assume that the implicit outputs of procedures it calls will remain unchanged if subsequently used as implicit inputs to those procedures or to companion procedures. (See Section 2.2.2.)
- All user code must be position independent. The data need not be position independent. However, for improved performance, data should be initialized to zero at compile or link time to avoid either position-independent constants or position-dependent addresses. (See Section 3.1.1.)
- Position-independent references (in a module) to another PSECT must use longword relative addressing so the VAX Linker can correctly allocate the data PSECT anywhere with respect to the code PSECT no matter how many code modules are included.
- External references must use general-mode addressing to allow the referenced procedures to be put in a shareable image without requiring changes to the calling program. (See Section 5.2.4.)
- Procedures cannot require their callers to pass dynamic string descriptors. (See Section 4.2.)
- Some procedure interface specifications retain state information from one call to the next, even though the procedures are not resource allocating. The interface specification uses one of the following techniques (in order of decreasing preference) to permit sequences of calls from independent parts of a program by either eliminating the use of static storage or overcoming its limitations:
 - 1 The interface specification consists of a sequence of calls to a set of one or more procedures — the first procedure allocates and returns (as an output argument to the calling program) one of the following:
 - The address of heap storage
 - Some other processwide identifying valueThis argument is passed to the other procedures explicitly by the calling program, and the last procedure deallocates any heap storage or processwide identifying value.
 - 2 The procedure's caller allocates all storage and passes the address on each call.

VMS Modular Programming Standard

A.3 Coding Rules

- 3 The interface specification consists of a single call, where the calling program passes the address of one or more action routines and arguments to be passed to them. The procedure calls the action routines during its execution. Results are retained by the procedure across calls to the action routines. (No static storage used.)
- 4 The interface specification consists of a sequence of calls to a set of one or more procedures. The first procedure, among other things, saves the contents of any still active static storage on a push down stack in heap storage, and the last procedure, among other things, restores the old contents of static storage. Thus, static storage is made available for implicit arguments to be passed from one procedure to the next in the sequence of calls (unknown to the calling program). However, if an exception can occur anywhere in the sequence, the calling program must establish a condition handler that calls the last procedure in the event of a stack unwind (to restore the old contents of static storage).

A.3.2 Initialization

- If a procedure requires initialization once for each image activation, it is done without the caller's knowledge by one of the following:
 - 1 Initializing at compile time
 - 2 Initializing at link time
 - 3 Adding a dispatch address to PSECT LIB\$INITIALIZE
 - 4 Testing and setting a statically allocated first-time flag on each call
- A procedure must not use LIB\$INITIALIZE to establish a condition handler before the main program is called if its action might conflict with that of other condition handlers established before the main program. For more information about initializing modular procedures, see Section 3.2.

A.3.3 Reporting Exception Conditions

- A procedure must not print error or informational messages either directly or by calling the \$PUTMSG system service. It must either return a condition value in R0 as a function value or call LIB\$SIGNAL or LIB\$STOP to output all messages. (LIB\$SIGNAL and LIB\$STOP may be called either directly or indirectly.) (See Section 2.4.)

A.3.4 AST Reentrancy

- To be AST-reentrant, a procedure must execute correctly while allowing any procedure (including itself) to be called between any two instructions. The other procedure can be an AST-level procedure, a condition handler, or another AST-reentrant procedure. (See Section 3.3.)
- A procedure that uses no static storage and calls only AST-reentrant procedures is automatically AST-reentrant. (See Section 3.3.3.)
- If a procedure uses static storage, it must use one of the following methods to be called from AST and non-AST levels:
 - Perform access and modification of the data base in a single uninterruptable instruction. This can be done only from VAX MACRO, and emulated instructions are not allowed. (See Section 3.3.4.1.)
 - Detect concurrency of data base access with "test and set" instructions at each access of the data base. (See Section 3.3.4.2.)
 - Keep a call-in-progress count incremented upon entry to the procedure and decremented upon return. (See Section 3.3.4.3.)
 - Disable AST interrupts on entry to the procedure and restore the state of the AST enables on return. (See Section 3.3.4.4.)
- If a procedure performs I/O from the AST level by calling VMS RMS \$GET and \$PUT system services, it must check for the record stream active error status (RMS\$_RSA). If this error is encountered, the procedure issues the \$WAIT system service and then retries the \$GET or \$PUT system service. (See Section 3.3.5.)
- A procedure should not depend on AST interrupts being disabled to execute correctly if there are other coding methods available. For example, RMS completion routines are implemented via ASTs and will not work if ASTs are disabled. (See Section 3.3.)

A.3.5 Resource Allocation

- A procedure should not allocate static storage unless it is a processwide, resource-allocating procedure or unless it must retain results for implicit inputs on subsequent invocations.
- Timing procedures and resource allocation procedures should make statistics available for performance evaluation and debugging by providing the entry points fac_SHOW_name and fac_STAT_name. (Optional. See Section 4.3.)
- If a procedure uses a processwide resource, it calls the appropriate resource allocating library procedure or system service to allocate the resource to avoid conflict with allocations made to other procedures. To conserve resources, a procedure that requests resource allocation does one of the following:
 - Calls the deallocation procedure before returning to the calling program

VMS Modular Programming Standard

A.3 Coding Rules

- Remembers the allocation in static storage and calls the deallocation procedure later
- Passes the responsibility for deallocation back to the calling program
- Allocates a fixed number of the resources independent of the number of times it is called

(See Section 2.4 and Section 3.1.5.)

A.3.6 The Format and Content of Coded Modules

- Each module must be documented with a module description. (See Section 2.5.1.)
- Each procedure must be documented with a procedure description. (See Section 2.5.2.)
- When symbol definitions are to be coordinated between more than one module, (such as control blocks, procedure argument values, and completion status codes), the definitions should be centralized in a common source file. Note, however, that the modules must be written in the same language. (See Section 3.1.3.)
- Instructions and statements should be uppercase, while comments are in upper and lowercase. (Optional. See Section 3.1.4.2.)
- Spaces should be added to improve readability. (Optional. See Section 3.1.4.3.)
- Block comments should be used. (Optional. See Section 3.1.4.4.)
- Use symbols rather than numbers in the body of the procedure. (Optional. See Section 3.1.4.1.)
- Procedure entry point names, module names, and PSECT names must conform to the naming conventions. (See Section 3.1.2.2, Section 3.1.2.4, and Section 3.1.2.5.)
- DIGITAL recommends that you also adhere to the naming conventions in choosing names for facilities and files. (Optional. See Section 3.1.2.1 and Section 3.1.2.3.)

A.3.7 Shareable Images

- A procedure's code is position independent. All references to external locations such as VMS System Service entry points, use general addressing mode. (Optional. See Section 5.2.4.)

A.3.8 Upward Compatibility

- When a new version of a procedure replaces an existing library procedure, all new arguments should be added at the end of the call sequence and made optional to maintain upward compatibility. (Optional. See Section 2.2.5 and Section 6.1.)
- A procedure's entry points are vectored using a separate MACRO module containing TRANSFER declarations. (Optional. See Section 5.2.2.)
- A procedure's code and data is position independent. (See Section 3.1.1.)

VMS Absolute Programming Standard

July 1984

Standard Form 100-100

1. The purpose of this standard is to provide a common basis for the development and testing of VMS Absolute programming. It is intended to be used by all users of the VMS Absolute programming language.

2. This standard is based on the VMS Absolute programming language as defined in the VMS Absolute programming language specification.

3. The standard is intended to be used by all users of the VMS Absolute programming language.

B Argument Characteristics

Each explicit argument is defined in terms of VMS usage, data type, access mechanism and passing mechanism. These four argument attributes are described in the sections below.

B.1 VMS Usage

The VMS usage entry indicates the abstract data structure of the argument. Table B-1 contains a list of the VMS data structures.

Table B-1 VMS Data Structures

Data Structure	Definition
access_bit_names	Homogeneous array of 32 quadword descriptors; each descriptor points to the name of one of the 32 bits in an access mask.
access_mode	Unsigned byte denoting a hardware access mode. This unsigned byte can take four values: 0 specifies kernel mode; 1, executive mode; 2, supervisor mode; and 3, user mode.
address	Unsigned longword denoting the virtual memory address of either data or code, but not of a procedure entry mask (which is of type "procedure").
address_range	Unsigned quadword denoting a range of virtual addresses, which identify an area of memory. The first longword specifies the beginning address in the range; the second longword specifies the ending address in the range.
arg_list	Procedure argument list consisting of one or more longwords. The first longword contains an unsigned integer count of the number of successive, contiguous longwords, each of which is an argument to be passed to a procedure by means of a VAX CALL instruction.
ast_procedure	Unsigned longword integer denoting the entry mask to a procedure to be called at AST level. (Procedures that are not to be called at AST level are of type "procedure".)
boolean	Unsigned longword denoting a boolean truth value flag. This longword may have only two values: 1 (true) and 0 (false).
byte_signed	This VMS data type is the same as the data type "byte (signed)" in Table B-2.

Argument Characteristics

B.1 VMS Usage

Table B-1 (Cont.) VMS Data Structures

Data Structure	Definition
byte_unsigned	This VMS data type is the same as the data type "byte integer (unsigned)" in Table B-2.
channel	Unsigned word integer that is an index to an I/O channel.
char_string	String of from 0 to 65,535 8-bit characters. This VMS data type is the same as the data type "character string" in Table B-2.
complex_number	One of the VAX standard complex floating-point data types.
cond_value	Unsigned longword integer denoting a condition value (that is, a return status or system condition code), which is typically returned by a procedure in R0.
context	Unsigned longword that is used by a called procedure to maintain position over an iterative sequence of calls. It is usually initialized by the caller, but thereafter manipulated by the called procedure.
date_time	64-bit unsigned, binary integer denoting a date and time as the number of elapsed 100-nanosecond units since 00:00 o'clock, November 17, 1858. This VMS data type is the same as the data type "absolute date and time" in Table B-2.
device_name	Character string denoting the name of a device. It can be a logical name, but if it is, it must translate to a valid device name.
ef_cluster_name	Character string denoting the name of an event flag cluster. It can be a logical name, but if it is, it must translate to a valid event flag cluster name.
ef_number	Unsigned longword integer denoting the number of an event flag.
exit_handler_block	Variable-length structure denoting an exit handler control block.
fab	Structure denoting an RMS file access block.

Argument Characteristics

B.1 VMS Usage

Table B-1 (Cont.) VMS Data Structures

Data Structure	Definition
file_protection	Unsigned word that is a 16-bit mask that specifies file protection. The mask contains four 4-bit fields, each of which specifies the protection to be applied to file access attempts by one of the four categories of user: from the rightmost field to the leftmost field, (1) system users, (2) the file owner, (3) users in the same UIC group as the owner, and (4) all other users (the world). Each field specifies, from the rightmost bit to the leftmost bit: (1) delete access, (2) execute access, (3) write access, (4) read access. Set bits indicate that access is denied.
floating_point	One of the VAX standard floating-point data types.
function_code	Unsigned longword specifying the exact operations a procedure is to perform. This longword has two word-length fields: the first field is a number specifying the major operation; the second field is a mask or bitvector specifying various suboperations within the major operation.
io_status_block	Quadword structure containing information returned by a procedure that completes asynchronously. The information returned varies depending on the procedure.
item_list_2	Structure that consists of one or more item descriptors and that is terminated by a longword containing 0. Each item descriptor is a 2-longword structure that contains three fields.
item_list_3	Structure that consists of one or more item descriptors and that is terminated by a longword containing 0. Each item descriptor is a 3-longword structure that contains four fields.
item_quota_list	Structure that consists of one or more quota descriptors and that is terminated by a byte containing a value defined by the symbolic name PQL\$_LISTEND. Each quota descriptor consists of a 1-byte quota name followed by an unsigned longword containing the value for that quota.
lock_id	Unsigned longword integer denoting a lock identifier. This lock identifier is assigned by the lock manager facility to a lock when the lock is granted.

Argument Characteristics

B.1 VMS Usage

Table B-1 (Cont.) VMS Data Structures

Data Structure	Definition
lock_status_block	Structure into which the lock manager facility writes status information about a lock. A lock status block always contains at least two longwords: the first word of the first longword contains a status code; the second word of the first longword is reserved to DIGITAL; and the second longword contains the lock identifier (VMS type "lock_id".)
lock_value_block	16-byte block that the lock manager facility includes in a lock status block if the user requests it. The contents of the lock value block are user-defined and are not interpreted by the lock manager facility.
logical_name	Character string of from 1 to 255 characters that identifies a logical name or equivalence name to be manipulated by VMS logical name system services. Logical names that denote specific VMS objects have their own VMS types: for example, a logical name identifying a device has the VMS type "device_name".
longword_signed	This VMS data type is the same as the data type "longword integer (signed)" in Table B-2.
longword_unsigned	This VMS data type is the same as the data type "longword (unsigned)" in Table B-2.
mask_byte	Unsigned byte wherein each bit is interpreted by the called procedure. A mask is also referred to as a set of "flags" or as a "bitmask".
mask_longword	Unsigned longword wherein each bit is interpreted by the called procedure. A mask is also referred to as a set of "flags" or as a "bitmask".
mask_quadword	Unsigned quadword wherein each bit is interpreted by the called procedure. A mask is also referred to as a set of "flags" or as a "bitmask".
mask_word	Unsigned word wherein each bit is interpreted by the called procedure. A mask is also referred to as a set of "flags" or as a "bitmask".
null_arg	Unsigned longword denoting a "null argument." A "null argument" is an argument whose only purpose is to "hold a place" in the argument list.
octaword_signed	This VMS data type is the same as the data type "octaword integer (signed)" in Table B-2.
octaword_unsigned	This VMS data type is the same as the data type "octaword (unsigned)" in Table B-2.

Argument Characteristics

B.1 VMS Usage

Table B-1 (Cont.) VMS Data Structures

Data Structure	Definition
page_protection	Unsigned longword specifying page protection to be applied by the VAX hardware. Protection values are specified using bits 0 to 3; bits 4 to 31 are ignored.
procedure	Unsigned longword denoting the entry mask to a procedure that is not to be called at AST level. (Arguments specifying procedures to be called at AST level have the VMS type "ast_procedure".)
process_id	Unsigned longword integer denoting a process identifier (PID). This process identifier is assigned by VMS to a process when the process is created.
process_name	Character string that specifies the name of a process.
quadword_signed	This VMS data type is the same as the data type "quadword integer (signed)" Table B-2.
quadword_unsigned	This VMS data type is the same as the data type "quadword (unsigned)" in Table B-2.
rights_holder	Unsigned quadword specifying a user's access rights to a system object. This quadword consists of two fields: the first is an unsigned longword identifier (VMS type "rights_id") and the second is a longword bitmask wherein each bit specifies an access right.
rights_id	Unsigned longword denoting a rights identifier, which identifies an interest group in the context of the VMS security environment. This rights environment may consist of all or part of a user's User Identification Code (UIC).
rab	Structure denoting an RMS record access block.
section_id	Unsigned quadword denoting a global section identifier. This identifier specifies the version of a global section and the criteria to be used in matching that global section.
section_name	Character string denoting a global section name. This character string can be a logical name, but it must translate to a valid global section name.
system_access_id	Unsigned quadword that denotes a system identification value that is to be associated with a rights database.
time_name	Character string specifying a time value in VMS format.
uic	Unsigned longword denoting a User Identification Code (UIC).

Argument Characteristics

B.1 VMS Usage

Table B-1 (Cont.) VMS Data Structures

Data Structure	Definition
user_arg	Unsigned longword denoting a user-defined argument. This longword is passed to a procedure as an argument, but the contents of the longword are defined and interpreted by the user.
varying_arg	Unsigned longword denoting a variable argument. A variable argument can have variable types, depending on specifications made for other arguments in the call.
vector_byte_signed	A homogeneous array whose elements are all signed bytes.
vector_byte_unsigned	A homogeneous array whose elements are all unsigned bytes.
vector_longword_signed	A homogeneous array whose elements are all signed longwords.
vector_longword_unsigned	A homogeneous array whose elements are all unsigned longwords.
vector_quadword_signed	A homogeneous array whose elements are all signed quadwords.
vector_quadword_unsigned	A homogeneous array whose elements are all unsigned quadwords.
vector_word_signed	A homogeneous array whose elements are all signed words.
vector_word_unsigned	A homogeneous array whose elements are all unsigned words.
word_signed	This VMS data type is the same as the data type "word integer (signed)" in Table B-2.
word_unsigned	This VMS data type is the same as the data type "word (unsigned)" in Table B-2.

B.2 Data Type

An argument's data type indicates the VAX data type that must be used for the argument.

Table B-2 contains the data types allowed by the VAX Procedure Calling and Condition Handling Standard.

Table B-2 VAX Standard Data Types

Data Type	Symbolic Code
Absolute date and time	DSC\$K_DTYPE_ADT
Byte integer (signed)	DSC\$K_DTYPE_B
Bound label value	DSC\$K_DTYPE_BLV
Bound procedure value	DSC\$K_DTYPE_BPV
Byte (unsigned)	DSC\$K_DTYPE_BU
COBOL intermediate temporary	DSC\$K_DTYPE_CIT
D_floating	DSC\$K_DTYPE_D
D_floating complex	DSC\$K_DTYPE_DC
Descriptor	DSC\$K_DTYPE_DSC
F_floating	DSC\$K_DTYPE_F
F_floating complex	DSC\$K_DTYPE_FC
G_floating	DSC\$K_DTYPE_G
G_floating complex	DSC\$K_DTYPE_GC
H_floating	DSC\$K_DTYPE_H
H_floating complex	DSC\$K_DTYPE_HC
Longword integer (signed)	DSC\$K_DTYPE_L
Longword (unsigned)	DSC\$K_DTYPE_LU
Numeric string, left separate sign	DSC\$K_DTYPE_NL
Numeric string, left overpunched sign	DSC\$K_DTYPE_NLO
Numeric string, right separate sign	DSC\$K_DTYPE_NR
Numeric string, right overpunched sign	DSC\$K_DTYPE_NRO
Numeric string, unsigned	DSC\$K_DTYPE_NU
Numeric string, zoned sign	DSC\$K_DTYPE_NZ
Octaword integer (signed)	DSC\$K_DTYPE_O
Octaword (unsigned)	DSC\$K_DTYPE_OU
Packed decimal string	DSC\$K_DTYPE_P
Quadword integer (signed)	DSC\$K_DTYPE_Q
Quadword (unsigned)	DSC\$K_DTYPE_QU
Character string	DSC\$K_DTYPE_T
Aligned bit string	DSC\$K_DTYPE_V
Varying character string	DSC\$K_DTYPE_VT
Unaligned bit string	DSC\$K_DTYPE_VU
Word integer (signed)	DSC\$K_DTYPE_W
Word (unsigned)	DSC\$K_DTYPE_WU
Unspecified	DSC\$K_DTYPE_Z
Procedure entry mask	DSC\$K_DTYPE_ZEM
Sequence of instruction	DSC\$K_DTYPE_ZI

Argument Characteristics

B.3 Access Mechanism

B.3 Access Mechanism

The argument access entry describes the way in which the called routine accesses the data specified by the argument. The following three methods of access are the most common (the formal argument is the procedure argument from the point of view of the called procedure, and the actual argument is the procedure argument from the point of view of the procedure's caller):

- Read only. The formal argument is a constant. The associated actual argument may only be read; it may not be written to or modified.
- Write only. The formal argument is a variable. The value of the associated actual argument is written into the variable by the called procedure.
- Modify. The formal argument is a variable. The value of this variable is written by the procedure's caller and may be read or modified by the called procedure.

The following is a complete list of the access types allowed by the VAX Procedure Calling and Condition Handling Standard:

- Read only
- Write only
- Modify
- Function call (before return)
- JMP after unwind
- Call after stack unwind
- Call without stack unwind

B.4 Passing Mechanisms

The argument passing mechanism is the way in which an argument specifies the actual data to be used by the called routine. There are three types of passing mechanisms:

- By value. When the longword argument in the argument list contains the actual data to be used by the routine, the actual data is said to be passed to the routine by value. Note that, since an argument is only one longword in length, only data that can be represented in one longword can be passed by value.
- By reference. When the longword argument in the argument list contains the address of the data to be used by the routine, the data is said to be passed by reference.
- By descriptor. When the longword argument in the argument list contains the address of a descriptor, the data is said to be passed by descriptor. A descriptor consists of two or more longwords (depending on the class of descriptor used), which describe the location, length, and data type of the data to be used by the called routine.

Figure B-1 illustrates the three passing mechanisms.

Table B-3 contains the passing mechanisms allowed by the VAX Procedure Calling and Condition Handling Standard.

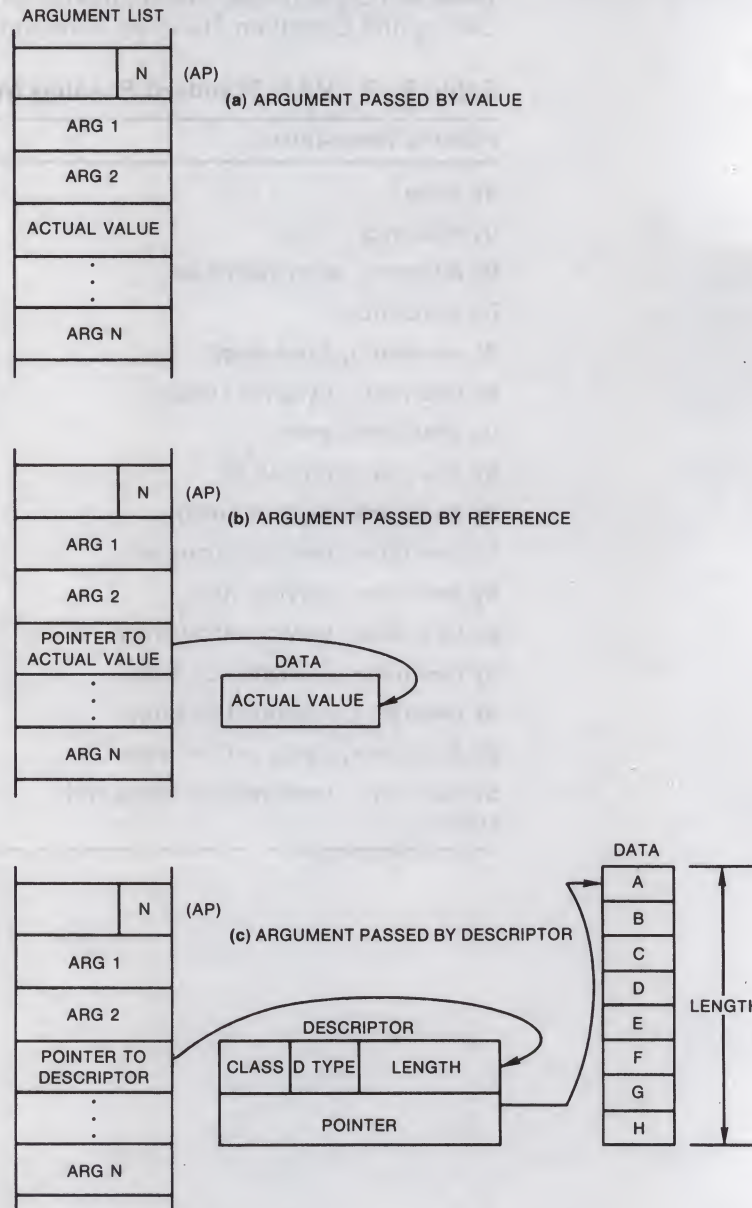
Table B-3 VAX Standard Passing Mechanisms

Passing Mechanism	Descriptor Code
By value	
By reference	
By reference, array reference	
By descriptor	
By descriptor, fixed-length	DSC\$K_CLASS_S
By descriptor, dynamic string	DSC\$K_CLASS_D
By descriptor, array	DSC\$K_CLASS_A
By descriptor, procedure	DSC\$K_CLASS_P
By descriptor, decimal string	DSC\$K_CLASS_SD
By descriptor, noncontiguous array	DSC\$K_CLASS_NCA
By descriptor, varying string	DSC\$K_CLASS_VS
By descriptor, varying string array	DSC\$K_CLASS_VSA
By descriptor, unaligned bit string	DSC\$K_CLASS_UBS
By descriptor, unaligned bit array	DSC\$K_CLASS_UBA
By descriptor, string with bounds	DSC\$K_CLASS_SB
By descriptor, unaligned bit string with bounds	DSC\$K_CLASS_UBSB

Argument Characteristics

B.4 Passing Mechanisms

Figure B-1 Procedure Argument Passing Mechanisms



Note: ARG 1, ARG 2, and ARG N can be passed by value, by reference, or by descriptor in any of these examples.

:(AP) = argument pointer

N = number of arguments

ZK-1962-84

Index

A

Argument

- access mechanism • B-8
- adding new • 6-3
- characteristics • B-1
- explicit • 2-3
- implicit • 2-3
- optional • 2-11, A-3
- order • 2-11, A-2
- passing mechanism • B-8
- VMS data types • B-6
- VMS usage • B-1

Argument blocks • 6-4

AST (asynchronous system trap)

- condition handling at AST level • 3-26
- definition • 3-19
- disabling interrupts • 3-24
- handler • 3-19, 3-21
- I/O at AST-level • 3-25, A-5
- interrupt • 3-19
- reentrancy • 3-19, 3-20, A-5
- routine • 3-19
- thread • 3-19
- writing AST-reentrant procedures • 3-20

Asynchronous system trap

- See AST

B

Black box testing • 4-2

Bound procedure value • 3-12

Busy wait • 3-21

C

Call-in-progress count • 3-24

Case

- using upper and lower • A-6

CMS (Code Management System) • 1-12

Code

- AST-reentrant • 3-19

Code (cont'd.)

- fully-reentrant • 3-19
- maintaining readability • 3-7
- position-independent • 3-1
- writing AST-reentrant procedures • 3-20

Coding guidelines • 3-1

Comment

- block • 3-9, A-6
- delimiters • 3-9

Common source files • 3-7, A-6

- declarations • 3-7

Condition handling

- at AST level • 3-26

Condition values • 3-3

D

Data types • B-6

Deadlock • 3-21

Design stage • 2-1

Documentation

- module description • 2-19, A-6
- procedure description • 2-20, A-6

DSC\$K_DTYPE_BP_V • 3-12

- See also User-action routine

DSC\$K_DTYPE_ZEM • 3-11

- See also User-action routine

E

Entry points

- See JSB entry points

Event flag • 2-16

F

Facility

- creation • 5-1
- library • 3-2
- naming • 5-1
- naming conventions • 3-2
- number • 3-3

Index

Facility (cont'd.)

prefix • 3-2, 5-1

First-time flag

testing and setting • 3-14

Full-reentrancy • 3-19

I

I/O • 2-16, A-4

asynchronous • 3-25

at AST-level • 3-25

file • 2-18

synchronous • 3-25

Initialization • 3-12, A-4

at run time • 3-17

of modular procedures • 3-12

of storage • 3-14

using LIB\$INITIALIZE • 3-17, A-4

Integration stage • 5-1

Integration testing • 4-1, 4-5

J

JSB entry points • 2-12, A-2

L

Language independence

testing for • 4-1, 4-4

Language-Sensitive Editor • 1-12

Levels of abstraction • 2-2

LIB\$INITIALIZE • 3-17

See also Initialization

Library

updating • 6-5

Library facility • 3-2

Linker options file

creating • 5-8

updating • 6-6

Lock manager • 3-21

Logical unit numbers • 2-16

M

MMS (Module Management System) • 1-12

Monitoring procedures • 4-8, A-5

in the Run-Time Library • 4-9

timer • 4-8

N

Naming conventions • 3-1, A-6

for facilities • 3-2

for files • 3-4

for modules • 3-4

for procedures • 3-3

for PSECTs • 3-5

O

Object module library

creating • 5-2

updating • 6-5

Organizing

files and modules • 2-1

procedures • 2-1

P

PCA (Performance and Test Coverage Analyzer) • 1-12

Performance analysis • 4-8

Position-independence • 3-1, A-3

Procedures

entry mask • 3-11

entry point names • 3-3

grouping • 5-1

interface • 2-3, A-2

libraries • 5-1

Program section

See PSECT

PSECT • 2-13, 3-5, A-3

DIGITAL-written • 3-5

LIB\$INITIALIZE • 3-17

user-written • 3-5

R

Race condition
 avoiding at AST-level • 3-21
 elimination of • 3-21
 Reentrancy
 AST • 3-19
 full • 3-19
 Regression testing • 6-1
 Returning condition values • 2-23
 RMS (Record Management Services) • 1-11
 Run-Time Library procedures • 1-6

S

SCA (Source Code Analyzer) • 1-13
 Screen management resources • 2-17
 Shareable image • A-6
 creating • 5-4
 updating • 6-6
 Shareable image library
 creating • 5-10
 updating • 6-7
 SHOW entry point • 4-8
 Signaling and condition handling • 2-22
 Signaling error conditions • 2-23
 Single instruction access • 3-22
 Software life cycle • 1-1
 STAT entry point • 4-9
 Storage • 2-12
 heap • 2-12
 initializing • 3-14
 stack • 2-12
 static • 2-13, A-5
 summary • 2-15
 Symbol
 cross-reference listing • 3-8
 definitions • A-6
 in place of numbers • 3-8, A-6
 System resources • 2-12
 System service • 3-11, A-2
 what is available • 1-8

T

Terminal I/O • 2-17
 Test and set instructions • 3-23
 Testing new procedures • 4-1
 black box • 4-2
 integration • 4-1, 4-5
 language independence • 4-1, 4-4
 modularity • 4-1
 reentrancy • 4-6
 regression • 6-1
 unit • 4-1
 white box • 4-3
 Threads of execution • 3-19
 Tools to aid in application development • 1-12
 Transfer vector
 changing • 6-6
 creating • 5-5
 updating • 6-3

U

Unit testing • 4-1
 black box • 4-2
 white box • 4-3
 Upward compatibility • 6-1, A-7
 User-action routine • 2-7
 interface • 3-11
 optional • 3-11
 passing • 3-11
 Using procedure libraries • 5-11
 Utility routines • 1-10

V

VAX DEC/CMS (Code Management System) • 1-12
 VAX DEC/MMS (Module Management System) • 1-12
 VAX DEC/Test Manager • 1-12
 VMS usage • B-1

W

White box testing • 4-3

8370W **NOTES**

230001 **NOTES**

NOTES

NOTES

NOTES

Reader's Comments

Guide to Creating VMS
Modular Procedures
AA-LA58A-TE

Please use this postage-paid form to comment on this manual. If you require a written reply to a software problem and are eligible to receive one under Software Performance Report (SPR) service, submit your comments on an SPR form.

Thank you for your assistance.

I rate this manual's:

	Excellent	Good	Fair	Poor
Accuracy (software works as manual says)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Completeness (enough information)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Clarity (easy to understand)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Organization (structure of subject matter)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Figures (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Examples (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Index (ability to find topic)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Page layout (easy to find information)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

I would like to see more/less _____

What I like best about this manual is _____

What I like least about this manual is _____

I found the following errors in this manual:

Page	Description
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____

Additional comments or suggestions to improve this manual:

I am using **Version** _____ of the software this manual describes.

Name/Title _____ Dept. _____

Company _____ Date _____

Mailing Address _____

_____ Phone _____

Do Not Tear - Fold Here and Tape

digitalTM



No Postage
Necessary
if Mailed
in the
United States



BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 33 MAYNARD MASS.

POSTAGE WILL BE PAID BY ADDRESSEE

DIGITAL EQUIPMENT CORPORATION
Corporate User Publications—Spit Brook
ZK01-3/J35 110 SPIT BROOK ROAD
NASHUA, NH 03062-9987



Do Not Tear - Fold Here

Cut Along Dotted Line

Reader's Comments

Guide to Creating VMS
Modular Procedures
AA-LA58A-TE

Please use this postage-paid form to comment on this manual. If you require a written reply to a software problem and are eligible to receive one under Software Performance Report (SPR) service, submit your comments on an SPR form.

Thank you for your assistance.

I rate this manual's:

	Excellent	Good	Fair	Poor
Accuracy (software works as manual says)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Completeness (enough information)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Clarity (easy to understand)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Organization (structure of subject matter)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Figures (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Examples (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Index (ability to find topic)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Page layout (easy to find information)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

I would like to see more/less _____

What I like best about this manual is _____

What I like least about this manual is _____

I found the following errors in this manual:

Page	Description
------	-------------

_____	_____
-------	-------

_____	_____
-------	-------

_____	_____
-------	-------

_____	_____
-------	-------

Additional comments or suggestions to improve this manual:

I am using **Version** _____ of the software this manual describes.

Name/Title _____ Dept. _____

Company _____ Date _____

Mailing Address _____

_____ Phone _____

Do Not Tear - Fold Here and Tape

digital™




No Postage
Necessary
if Mailed
in the
United States

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 33 MAYNARD MASS.

POSTAGE WILL BE PAID BY ADDRESSEE

DIGITAL EQUIPMENT CORPORATION
Corporate User Publications—Spit Brook
ZK01-3/J35 110 SPIT BROOK ROAD
NASHUA, NH 03062-9987



Do Not Tear - Fold Here

Cut Along Dotted Line

digital